arXiv:1612.07811v1 [stat.ME] 22 Dec 2016

Bootstrap inference after using multiple queries for model selection

Jelena Markovic, Jonathan Taylor

Department of Statistics
Stanford University
Sequoia Hall
Stanford, CA 94305, USA
e-mail: jelenam@stanford.edu
jonathan.taylor@stanford.edu

Abstract: Recently, [Tian and Taylor, 2015b] developed a selective inference approach with a randomized response, demonstrating that such randomization provides a way to tradeoff the information used for model selection and the *leftover* information, used for inference about model parameters. They constructed an asymptotically pivotal test statistic using a CLT that holds *without selection*. Adjusting the resulting Gaussian limit for selection via a *selective CLT* allows for selective inference in non-parametric settings.

In this work, we provide a refinement of their selective CLT result, most notably we relax their local alternatives assumption. Under some regularity assumptions on the density of the randomization, including heavier tails than Gaussian satisfied by e.g. logistic distribution, we prove the selective CLT holds without any assumptions on the underlying parameter, allowing for rare selection events. We also show that under the local alternatives assumption on the parameter, selective CLT holds for Gaussian randomization as well, though the quantitative results are qualitatively different for the Gaussian randomization as compared to the heavier tailed results.

Furthermore, we propose a bootstrap version of this test statistic (bootstrapped pivot). We prove that the bootstrap test statistic is also asymptotically pivotal uniformly across a family of non-parametric distributions. This result can be interpreted as resolving the impossibility results of [Leeb and Pötscher, 2006a]. We describe a way of using the wild bootstrap and projected Langevin Monte Carlo sampling method to compute a bootstrapped test statistic and the corresponding confidence intervals valid after selection.

As most data analysts will want to try various model selection algorithms when choosing a model, we present a way to construct confidence intervals after multiple views/queries of the data. In this setting, an analyst runs several model selection procedures on the same data and choses a parameter of interest upon seeing the outcomes of all of the model selection procedures. We note that this target of interest need not agree exactly with the results of any model selection procedure – the data analyst can use their own expertise to choose a final parameter of interest but is allowed access to the results of the model selection procedure before choosing their parameter of interest. We construct the selective confidence intervals for the selected parameter using both pivot constructions, plugin CLT and the bootstrap.

Finally, we compare our methods to data splitting, in which some portion of the data is used for model selection (stage 1) and the remaining data used for inference (stage 2). In [Fithian et al., 2014], it was noted that one can often improve on data splitting by using information leftover after stage 1, though the examples and results were in a parametric setting. In this work, we construct both inferential and sampling tools to reuse the information in the data from the first stage and provide tests with greater power than traditional data splitting.

All the computations can be done with any of the examples from [Tian Harris et al., 2016], including GLMs, forward-stepwise and marginal screening, and their combination into multiple views framework. We present the implementation results on some of these.

1. Introduction

This work continues a line of research in inference after model selection beginning with [Lee et al., 2016]. In this paper, the authors use the LASSO to select variables and provide valid tests and confidence intervals conditional on the result of the selection by the LASSO. Of course, these tests are valid

only for the statistical model for which they were designed. In [Lee et al., 2016], the model was the saturated model with Gaussian errors of known variance proposed by [Berk et al., 2013]. Followup work, including [Fithian et al., 2014, Tian and Taylor, 2015b] among others, point out that this model may not be realistic in many situations. However, the principle of conditioning on the result of the selection by the LASSO is applicable in other statistical models besides the saturated model. This train of thought leads to the concept of a *selective model* in which each distribution $\mathbb{F}$ in some statistical model $\mathcal{M}$ is conditioned on the output of some model selection procedure, a canonical example being the choice of variables and their signs by the LASSO. The selective model is determined by the outcome M of the procedure and the corresponding selective likelihood ratio

$$\ell^*_{\mathbb{F},E}(\boldsymbol{S}) = \frac{d\mathbb{F}^*}{d\mathbb{F}}(\boldsymbol{S}) = \frac{\mathbb{P}_{\mathbb{F}}\{\widehat{M} = M | \boldsymbol{S}\}}{\mathbb{P}_{\mathbb{F}}\{\widehat{M} = M\}}$$

where $\widehat{M}$ is some possibly randomized selection procedure which we might think of as a *query*, i.e. a function we evaluate on data $\mathcal{S}$ (along with possible randomization) and are given the result. The behavior of the random variables

$$(\ell^*_{\mathbb{F},E})_{\mathbb{F}\in\mathcal{M}}$$

can be used to determine properties of the selective model

$$\mathcal{M}_E = \left\{ \mathbb{F}^* : \frac{d\mathbb{F}^*}{d\mathbb{F}} = \ell^*_{\mathbb{F},E}, \mathbb{F} \in \mathcal{M} \right\}.$$

For example, [Tian and Taylor, 2015a] shows that consistency under a sequence of *unselective* (or, pre-selective) models $\mathcal{M}_n$ can be transferred to a corresponding sequence of selective models $\mathcal{M}^*_{n,E_n}$ under conditions on the appropriate selective likelihood ratios. Similarly, questions of weak convergence for sequences of selective models can be related to weak convergence under the unselective models. It is these questions of weak convergence that are the focus of this work. In particular, we construct selective bootstrap procedures that produce asymptotically pivotal quantities which converge weakly in a uniform sense over a large class of models.

Before detailing our contributions here, we review other related work in selective inference. Assuming Gaussian errors, the truncated Gaussian test statistic constructed in [Lee et al., 2016] is an exact pivot valid after selection. It is constructed for saturated models, meaning that no relationship is assumed between the response vector and the predictors before or after selection. In the saturated model framework, the selection is done to adjust for the choice of the parameter to report. Some recent work in selective inference has been focusing on removing the Gaussian assumption on the data and getting asymptotic tests in non-parametric settings. [Tian and Taylor, 2015a] and [Tibshirani et al., 2015] prove the truncated Gaussian statistic of [Lee et al., 2016] is asymptotically pivotal in the unconditional sense, i.e. the convergence results are under the original distribution of the data before selection.[1] [Tibshirani et al., 2015] also proposes a bootstrap version of this test that is asymptotically conservative again in the unconditional sense. All the works mentioned have been focusing on the non-randomized selective inference, meaning that given the data there is no additional randomness coming in either the model selection step or the inferential procedures.

[Tian and Taylor, 2015b] propose doing selection on a randomized version of the response vector or adding randomization directly into the objective function of a model selection algorithm. They notice a significant increase in the power of a test with added randomization. Also, by adding heavy-tailed randomization (e.g. Laplace or logistic), the test statistic they construct behaves better than the truncated Gaussian statistic in the sense that the denominator of their test statistic is bounded away from zero. They prove the test statistic is asymptotically pivotal under the conditional distribution

[1] For more discussion on the unconditional perspective for selective inference see [Tibshirani et al., 2015].

of the data in the uniform sense across a non-parametric family of distributions.[2] Their work allows for both saturated and selected model framework. In the selected model framework, we assume some relationship between the response and the covariates after looking at the outcome of the model selection algorithm [Fithian et al., 2014]. In the selected model framework, inference is done by conditioning on less than in saturated (the difference being the null statistic), hence increasing the power. In this scenario, the conditioning on the selection region adjusts for both choosing the model and the choice of a parameter to test.[3] Here, we build on the main ideas of randomized selective inference.

1.1. Contributions of this paper

From theoretical perspective, we first revisit the selective CLT result of [Tian and Taylor, 2015b] and relax their assumptions. Under some regularity conditions on the randomization, we prove selective CLT without assuming local alternatives. This allows for the parameter vector to be arbitrarily far from the selection region, allowing for selection events that are relatively rare under the true data generating mechanism. In this sense, such procedures are robust to such rare events. This result requires the randomization have tails heavier than Gaussian such as logistic. However, Gaussian randomization occurs naturally in problems like data splitting. In this case, we use the data and the random split to construct the randomization that is asymptotically normal. This problem then becomes randomized inference problem with Gaussian randomization. Our second theoretical result is the proof that selective CLT holds with Gaussian randomization under the local alternatives. This means if the distance of the parameter vector to the selection region is not growing too fast, the constructed pivot is asymptotically uniform hence valid for inference. [4] As our results for Gaussian randomization require some form of local alternatives, they are less robust to rare selection events than corresponding heavier tailed randomizations. While we have tried to remove such assumptions for Gaussian randomization, our current method of proof requires such assumptions. However, in the non-randomized setting (which can be thought of as randomization with a degenerate distribution) local alternatives are necessary. See [Tian and Taylor, 2015b] for a univariate example in which the pivotal quantity fails to converge weakly.

Building on the non-parametric results, we also propose a bootstrap version of the asymptotically pivotal test statistic constructed in [Tian and Taylor, 2015b]. We prove that the bootstrap test statistic is asymptotically pivotal under the conditional distribution of the data and in uniform sense across a family of non-parametric distributions. In order to prove the consistency of bootstrap, we assume the asymptotic uniformity of the pivot (the result of the selective CLT holds without bootstrap), local alternatives and, additionally, we assume the consistency of the bootstrap without selection and that we are able to consistently estimate the variance. Under these mild conditions, the bootstrapped pivot is then also valid post-selection and can be used for inference.

When it comes to computing the bootstrap pivot, there are two important issues that we address here. The first one is sampling data from the density conditional on the selection region (called selective density) which introduces complicated constraints. To address this, we use the pull-back trick of [Tian et al., 2016] and [Tian Harris et al., 2016] to sample data and the optimization variables from a simpler region instead of sampling data and the randomization variables from a much more complicated region.[5] The second issue is using the bootstrapped samples for the data. Under

[2]Conditional distribution refers to the distribution of the data after selection, i.e. by conditioning on the selected model. Since the conditional statements imply unconditional ones, the guarantees under the conditional distribution are stronger.

[3]For further distinctions on selected and saturated model without randomization see [Fithian et al., 2014], which introduces the concept of the selected model.

[4]Even though there seem to be differences between Gaussian and logistic randomization in the assumptions needed for selective CLT to hold, empirically, as we will see later in the implementation results, both work fine in practice.

[5]The optimization variables along with the data describe the selection region.

the selective density, not all bootstrapped samples are equally likely so standard resampling techniques do not obviously work. To address this issue, we use the wild bootstrap of [Wu, 1986] and [Mammen, 1993] since it allows us to sample the weights along with the optimization variables from a continuous density with constraints.

We develop two samplers for computing both the plugin Gaussian (also called plugin CLT, the pivot introduced in [Tian and Taylor, 2015b]) and the bootstrap pivot. We use both samplers to do valid inference after running several model selection procedures on the data, i.e. after what we call multiple views or multiple queries of the data. This allows a statistician to do inference after GLMs with ℓ_1 penalties, forward-stepwise, marginal screening etc. or any of their combinations.

We should emphasize that in order to do valid inference we do not assume the selected model is true but we construct confidence intervals for the population parameters that are chosen upon looking at the selected model. Also, upon looking at the outcomes of the model selection procedures, an analyst decides on her own model. Furthermore, this work allows for testing both individual parameters (selective t-test) or any linear combination (selective F-test).

Another important contribution that comes as an application of this work is a novel way to do data splitting through *data carving*, [Fithian et al., 2014]. Classical data splitting uses a part of the data (stage 1) to select a model and the leftover data for inference. Hence, the classical data splitting conditions on the whole first stage data used for selection. In this paper, we select the model using the stage 1 data as well, but by conditioning only on the model selected using the first stage data, we use the whole data for inference. Conditioning on less, the data carving procedure provides an increase in power compared to the traditional data splitting, hence shorter confidence intervals. Technically, we rewrite the model selection problem using the first stage data as the randomized version of the model selection problem using the full data with the randomization being a function of both the first stage and the second stage data. This way the sampling problem becomes the same as the sampling problem from [Tian Harris et al., 2016], where in the data carving setting the randomization is Gaussian asymptotically independent of the data. Then we reuse the sampler from [Tian Harris et al., 2016] to get the samples from the selective density that are used for inference. We also use the bootstrap sampler proposed in this work to construct the test statistic valid in this setting.

Finally, the multiple views framework is combined with the data splitting framework into the multiple splits, a way to do inference after looking at the models selected from several splits. All the sampling tools used here are novel tools with inferential guarantees.

1.2. Outline

The outline of this paper is as follows. We provide refinements of the selective CLT result of [Tian and Taylor, 2015b] (Section 2). Most notably, we do not require the local alternatives assumption allowing for the true underlying parameter to be arbitrary far from the selection region. We show the estimator consistent under the original (pre-selection) distribution of the data is also consistent in the conditional distribution of the data (Section 3). This requires some regularity conditions on the randomization density, satisfied by the logistic distribution for example. We also prove a selective CLT in the case of Gaussian randomization under the local alternatives assumption, thus extending the existing results (Section 2.3).

In Section 4, we propose a bootstrap statistic and show it is asymptotically pivotal after randomized model selection procedures with affine constraints. Here, we do assume local alternatives, hence the distance of the underlying parameter to the selection region cannot grow too fast with the sample size.

In Section 5, we describe two concrete examples of computing the bootstrap test statistic using the wild bootstrap. Other examples from [Tian Harris et al., 2016] can be done in a similar fashion. The first example uses Lasso as a model selection procedure and the second one uses marginal screening. All the examples are done with a random design matrix.

Further, in Section 6, we present a general way of doing inference after multiple views of the data, where each view corresponds to running a randomized model selection procedure and looking at its outcome. We illustrate the sampling procedure after multiple views of the data using two examples: forward-stepwise (Section 6.1) and randomized ℓ_1-penalized logistic regressions (Section 6.2).

In Section 7, we describe an alternative to data splitting, a more powerful inferential tool called data carving [Fithian et al., 2014]. Combining this technique with the multiple views of the data, we have the tools to do inference after multiple splits (Section 7.2).

We advocate the projected Langevin method (which enjoys theoretical guarantees described in [Bubeck et al., 2015]) as a sampling tool for the bootstrap weights and the optimization variables from the log-concave density with simple constraints (Section 8).

1.3. General setup

We describe a general framework of computing the bootstrap pivot in randomized selective inference. We do selection by adding a linear randomization term in the objective of the model selection algorithm. We then compute the bootstrapped test statistic while adjusting for selection, hence accounting/conditioning on the fact that an analyst looked at the selected model in order to perform inference.

Given data $\boldsymbol{S} \in \mathbb{R}^{n\times p}$ with rows $\boldsymbol{S}_i \overset{i.i.d}{\sim} \mathbb{F}_n$, $i = 1, \ldots, n$, where the distribution $\mathbb{F}_n$ has a density f, we solve the randomized model selection algorithm, introduced in [Tian and Taylor, 2015b], [Tian Harris et al., 2016], of the following form

$$\underset{\boldsymbol{\beta}\in\mathbb{R}^p}{\text{minimize}}\; \ell(\boldsymbol{\beta}; \boldsymbol{S}) + \mathcal{P}(\boldsymbol{\beta}) - \boldsymbol{\omega}^T\boldsymbol{\beta} + \frac{\epsilon}{2}\|\boldsymbol{\beta}\|_2^2, \quad \boldsymbol{S} \times \boldsymbol{\omega} \sim \mathbb{F}_n \times \mathbb{G}, \tag{1}$$

where $\ell(\boldsymbol{\beta}; \boldsymbol{S})$ is the loss function, $\mathcal{P}(\boldsymbol{\beta})$ is a penalty term, $\boldsymbol{\omega}$ is a randomization random variable drawn from a known distribution $\mathbb{G}$ on $\mathbb{R}^p$ with density g and $\epsilon > 0$ is a known constant.[6] Denote the solution of (1) as $\widehat{\boldsymbol{\beta}}(\boldsymbol{S}, \boldsymbol{\omega})$ and the selected model as $\widehat{M}(\boldsymbol{S}, \boldsymbol{\omega}) = M$.

In order to compute the null distribution of any test statistic after selection, we need to sample $(\boldsymbol{S}, \boldsymbol{\omega})$ such that $\widehat{M}(\boldsymbol{S}, \boldsymbol{\omega}) = M$, meaning that running (1) on the new data and randomization pairs selects the same model. Using the pull-back measure trick of [Tian Harris et al., 2016], we do not need to sample $(\boldsymbol{S}, \boldsymbol{\omega})$ but the data $\boldsymbol{S}$ and optimization variables $\boldsymbol{v} \in \mathbb{R}^m$ (which are problem dependent).[7] The optimization variables are chosen such that we can recover $\boldsymbol{\omega}$ by the subgradient equation of (1) as $\boldsymbol{\omega} = \omega(\boldsymbol{S}, \boldsymbol{v})$ for some function ω. Also, the selected model usually can be described only based on the optimization variables $\boldsymbol{v}$, hence $M = \widetilde{M}(\boldsymbol{v})$. This implies that selecting model M is equivalent to having the optimization variables, vector $\boldsymbol{v}$, constrained to $\mathcal{V}_M = \{\boldsymbol{v}' : \widetilde{M}(\boldsymbol{v}') = M\} \subset \mathbb{R}^m$. In selective inference problems, this constraint set generally becomes simpler than $\{(\boldsymbol{S}, \boldsymbol{\omega}) : \widehat{M}(\boldsymbol{S}, \boldsymbol{\omega}) = M\}$, thus the re-parametrization of the data and randomization in terms of the data and the optimization variables is extremely useful for sampling purposes.

In order to compute the distribution of any test statistic under the null, we sample $(\boldsymbol{S}, \boldsymbol{v})$ from a density proportional to

$$\left(\prod_{i=1}^n f(\boldsymbol{S}_i)\right) \cdot g(\omega(\boldsymbol{S}, \boldsymbol{v})) \cdot |J(\boldsymbol{S}, \boldsymbol{v})|, \tag{2}$$

supported on $\mathbb{R}^{n\times p} \times \mathcal{V}_M$, where $J(\boldsymbol{S}, \boldsymbol{v})$ is the Jacobian coming from the change of density. This sampling requires knowing the distribution of the data $\mathbb{F}_n$ exactly or using its asymptotic distribution. We compute plugin CLT pivot based on the data samples from the density above.

[6] Adding a quadratic term $\frac{\epsilon}{2}\|\boldsymbol{\beta}\|_2$ ensures there exists a solution of (1).

[7] Later in the examples we will see that we do not need to sample the whole data matrix $\boldsymbol{S}$ but only a function $T(\boldsymbol{S})$ of it, usually of dimension p, since the optimization problem in (1) depends on $\boldsymbol{S}$ only through $T(\boldsymbol{S})$ and having $T(\boldsymbol{S})$ samples suffice for inference.

A novel idea of this paper is to use the bootstrap distribution of the data hence to sample $(\boldsymbol{S}^*, \boldsymbol{v})$ such that $\widehat{M}(\boldsymbol{S}^*, \omega(\boldsymbol{S}^*, \boldsymbol{v})) = M$, where $\boldsymbol{S}^*$ is the bootstrapped data. Based on these samples, we construct pivots and also confidence intervals valid post-selection.

Standard non-parametric bootstrap of Efron [Efron, 1979] resamples $\boldsymbol{S}_i^*$, $i = 1, \ldots, n$, from the empirical distribution of the rows of the data matrix. This is equivalent to sampling b_i from a multinomial with equal probabilities $\frac{1}{n}$ at each $(1, \ldots, n)$ and setting $\boldsymbol{S}_i^* = \boldsymbol{S}_{b_i}$ for each $i = 1, \ldots, n$. To do bootstrap after selection, i.e. after conditioning on a model M, we would need to sample $(\boldsymbol{b}, \boldsymbol{v}) \in \{1, \ldots, n\}^n \times \mathbb{R}^m$ such that the model selection algorithm for the bootstrapped data $\boldsymbol{S}^*(\boldsymbol{b})$ and $\omega(\boldsymbol{S}^*(\boldsymbol{b}), \boldsymbol{v})$ gives the same model M, where $\boldsymbol{S}^*(\boldsymbol{b})$ denotes the matrix with rows $\boldsymbol{S}_i^* = \boldsymbol{S}_{b_i}$, $i = 1, \ldots, n$. This couples $\boldsymbol{b}$ and $\boldsymbol{v}$ so that their joint is proportional to

$$g(\omega(\boldsymbol{S}^*(\boldsymbol{b}), \boldsymbol{v})) \cdot |J(\boldsymbol{S}^*(\boldsymbol{b}), \boldsymbol{v})|,$$

and supported on $\{1, \ldots, n\}^n \times \mathcal{V}_M$. Since sampling from this density is hard, we use continuous versions of bootstrap, e.g. wild bootstrap or Bayesian bootstrap, instead of Efron's bootstrap.[8][9]

Wild bootstrap approximates the distribution of $\boldsymbol{S}$ before selection with the distribution of $\boldsymbol{S}^*(\boldsymbol{\alpha}) = \text{diag}(\boldsymbol{\alpha})\boldsymbol{S}$, where $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_n)$ and $\text{diag}(\boldsymbol{\alpha})$ has the elements of $\boldsymbol{\alpha}$ on its diagonal and zeroes elsewhere. We take $\alpha_i \overset{i.i.d.}{\sim} \mathbb{F}_\alpha$ with density f_α. This allows us to sample $(\boldsymbol{\alpha}, \boldsymbol{v})$ from

$$\left(\prod_{i=1}^{n} f_\alpha(\alpha_i) \right) \cdot g(\omega(\boldsymbol{S}^*(\boldsymbol{\alpha}), \boldsymbol{v})) \cdot |J(\boldsymbol{S}^*(\boldsymbol{\alpha}), \boldsymbol{v})|, \tag{3}$$

supported on $(\text{supp}(\mathbb{F}_\alpha))^n \times \mathcal{V}_M$. Now we can use some of the Monte Carlo techniques to sample from the density above with constraints (see Section 8). Since we sample the weights ($\boldsymbol{\alpha} \in \mathbb{R}^n$ if we take $\mathbb{F}_\alpha$ to have the full support) and the optimization variables ($\boldsymbol{v} \in \mathcal{V}_M$, usually a simple set), the constrains will usually be simple in most of the problems of interest hence the sampling is efficient. In the problems we consider, the Jacobian is a constant hence we do not need to compute it but there are selective inference problems in which the Jacobian is non-trivial [Tian Harris et al., 2016]. Given that we have an explicit form for the density in (3), we can easily sample form it.

Remark 1 *For the wild bootstrap to be consistent, it suffices to have the mean of $\mathbb{F}_\alpha$ to be 0 and the variance to be 1. The condition that the skewness of $\mathbb{F}_\alpha$ is also 1 has been introduced by [Liu, 1988] to improve the rate of convergence of the bootstrap estimate. In practice (Section 8.1), we use $\mathbb{F}_\alpha$ to be standard normal.*

There are other possible continuous bootstrap versions that can be used here but we do not pursue them in practice. The Bayesian bootstrap of [Rubin et al., 1981] or, more generally, the weighted bootstrap of [Praestgaard and Wellner, 1993] approximates the distribution of $\boldsymbol{S}$ before selection with the distribution of $\boldsymbol{S}^*(\boldsymbol{w}) = \text{diag}(\boldsymbol{w})\boldsymbol{S}$, where $\boldsymbol{w} = (w_1, \ldots, w_n)$ comes from a distribution $\mathbb{F}_{\boldsymbol{w}}$ with density $f_{\boldsymbol{w}}(\boldsymbol{w})$ supported on a simplex $\Delta_n = \{w_i \in [0, 1] : \sum_{i=1}^n w_i = 1\}$.[10] An example of the distribution on $\boldsymbol{w}$ is the Dirichlet distribution [Rubin et al., 1981]. In this case, the joint density of $(\boldsymbol{w}, \boldsymbol{t})$ is

$$f_{\boldsymbol{w}}(\boldsymbol{w}) \cdot g(\omega(\boldsymbol{S}^*(\boldsymbol{w}), \boldsymbol{v})) \cdot |J(\boldsymbol{S}^*(\boldsymbol{w}), \boldsymbol{v})|,$$

supported on $\Delta_n \times \mathcal{V}_M$, hence slightly more complicated than wild bootstrap proposal.

[8] There are some possible modifications of Efron's bootstrap that sample b_i's from the Poisson distribution instead of multinomial [Hanley and MacGibbon, 2006]; however, using this bootstrap after selection requires sampling from the discrete distributions we do not pursue that direction here.

[9] Computing the selective pivot directly using standard bootstrap with replacement is theoretically possible and will be elaborated on in Section 4; however, it involves weighting bootstrap samples with selection probabilities that are hard to approximate.

[10] For more restrictions on the distribution $\mathbb{F}_{\boldsymbol{w}}$ see [Praestgaard and Wellner, 1993].

1.4. Simple example - inference for the mean after randomized selection

We present an example where we do inference for the mean of i.i.d. random variables after we selected only the data vectors for which the randomized t-statistic is above some given threshold. The data vector $\boldsymbol{Y} = (Y_1, \dots, Y_n)$ consists of $Y_i \overset{i.i.d.}{\sim} \mathbb{F}_n$, $1 \leqslant i \leqslant n$. We want to test for the mean $H_0 : \mu(\mathbb{F}_n) = \mu$ after selection of the following form

$$\sqrt{n}\bar{Y} + \omega > \bar{a}, \quad (\boldsymbol{Y}, \omega) \sim \mathbb{F}_n^n \times \mathbb{G},$$

where $\bar{Y} = \frac{1}{n}\sum_{i=1}^n Y_i$, $\bar{a}$ is some fixed (and known) threshold and ω is added randomization independent of the data, $\omega \sim \mathbb{G}$, where the distribution $\mathbb{G}$ on $\mathbb{R}$ is known. Also, denote with $\bar{G}$, the survival function of ω, i.e. $\bar{G}(x) = \mathbb{G}\{\omega > x\}$. Furthermore, denote the distribution of the data vector $\boldsymbol{Y}$ after selection as $\mathbb{F}_n^*$ (post-selection distribution of the data).

Let us denote as $F_\mu(t)$ the CDF of $\sqrt{n}(\bar{Y} - \mu)$, where $\boldsymbol{Y} \sim \mathbb{F}_n^*$. [Leeb and Pötscher, 2006a] show that for the normal distribution $\mathbb{F}_n = \Phi(\mu, 1)$ one cannot estimate $F_\mu(t)$ uniformly consistently, i.e. for a fixed t any estimate $\widehat{F}(t)$ that is consistent at $\mu = 0$,

$$\widehat{F}(t) - F_0(t) \overset{\mathbb{F}_n^*}{\to} 0 \quad \text{as } n \to \infty,$$

(convergence in probability over the conditional distribution of the data) necessarily satisfies

$$\liminf_{n\to\infty} \sup_{|\mu| < \frac{h}{\sqrt{n}}} \mathbb{F}_n^* \left\{ |\widehat{F}(t) - F_\mu(t)| > \delta \right\} \geqslant c > 0,$$

for some constants $h, \delta, c > 0$, where the local alternatives can be replaced by $|\mu| < R$ for some constant R as well.[11] Hence, no uniformly consistent estimator of the conditional distribution $F_\mu(t)$ exists, not even locally.

However, we can still get uniformly consistent confidence interval for the mean μ by inverting a test. If we construct a test statistic $\widehat{F}_\mu(t)$ such that

$$\lim_{n\to\infty} \sup_{|\mu| < \frac{h}{\sqrt{n}}} \mathbb{F}_n^* \left\{ \sup_{t\in\mathbb{R}} \left| \widehat{F}_\mu(t) - F_\mu(t) \right| > \delta \right\} = 0, \tag{4}$$

then the confidence interval constructed by inverting this test statistic,

$$\mathcal{I}_\alpha(\boldsymbol{Y}) = \left\{ \mu : \widehat{F}_\mu\left(\sqrt{n}(\bar{Y} - \mu)\right) \in \left[\frac{\alpha}{2}, 1 - \frac{\alpha}{2}\right] \right\},$$

is uniformly valid under the conditional distribution of the data, i.e.

$$\liminf_{n\to\infty} \inf_{|\mu| < \frac{h}{\sqrt{n}}} \mathbb{F}_n^*\{\mu \in \mathcal{I}_\alpha(\boldsymbol{Y})\} \geqslant 1 - \alpha,$$

for a given level α (follows from Lemma A.1 in [Romano et al., 2012]). Our goal is to construct $\widehat{F}_\mu(t)$ that satisfies (4) using bootstrap.

Remark 2 *The non-regularity problems of this kind (parameter at the boundary) have been also considered in [McKeague and Qian, 2015] and [Laber et al., 2014]. [McKeague and Qian, 2015] also provides a bootstrap test for testing the global null in the non-randomized version of the model selection problem (marginal screening as in Section 5.2). As pointed out in [Leeb, 2015], their test is for a point null, hence the non-regularity problem is not an issue anymore [Leeb and Pötscher, 2006a],*

[11] They show the result for non-randomized selection events but the argument holds with randomized selection events as well.

[Leeb and Pötscher, 2006b]. We provide a similar fix, i.e. circumventing the non-regularity issue by providing a test, instead of doing inference based on a cumulative distribution function of an estimator (that is known to provide non-uniform limiting behavior [Leeb and Pötscher, 2006b]). However, our test is not just for testing a global null but can be used more generally for testing any linear combination of the underlying parameter vector. As for the construction of confidence intervals, the test of [McKeague and Qian, 2015] is hard to invert, while we also construct confidence intervals as well (Section 8.1).

[Laber et al., 2014] addresses non-regularity problem in a specific example of dynamic treatments, providing bootstrap confidence intervals that are conservative.

In this example, we can use standard bootstrap of Efron (sampling with replacement) to approximate $F_\mu(t)$ using

$$\widehat{F}^b_\mu(t) = \widehat{\mathbb{F}}^n_n \times \mathbb{G}\left\{\sqrt{n}(\bar{Y}^* - \bar{Y}) \leqslant t \,\middle|\, \sqrt{n}(\bar{Y}^* - \bar{Y}) + \sqrt{n}\mu + \omega > \bar{a}\right\},$$

where $\widehat{\mathbb{F}}_n$ is the empirical distribution of our sample $\boldsymbol{Y}$. Here, $Y^* = (Y_1^*, \ldots, Y_n^*) \sim \widehat{\mathbb{F}}^n_n$ denotes a bootstrap sample and $\bar{Y}^*$ its mean.

Computationally, we bootstrap the mean $\bar{Y}$ and get the bootstrapped mean samples as $\bar{Y}^{*b}$, $b = 1, \ldots, B$. Then we compute $\widehat{F}^b_\mu(t)$ as

$$\frac{\sum_{b=1}^B \bar{G}\left(\bar{a} - \sqrt{n}\mu - \sqrt{n}(\bar{Y}^{*b} - \bar{Y})\right) \cdot \mathbb{I}_{\{\sqrt{n}(\bar{Y}^{*b} - \bar{Y}) \leqslant t\}}}{\sum_{b=1}^B \bar{G}\left(\bar{a} - \sqrt{n}\mu - \sqrt{n}(\bar{Y}^{*b} - \bar{Y})\right)}. \tag{5}$$

In this case, the test statistic becomes $\widehat{F}^b(\sqrt{n}(\bar{Y} - \mu))$ and the confidence intervals can be constructed by inverting this test which involves a grid search for μ such that $\widehat{F}_\mu(\sqrt{n}(\bar{Y} - \mu)) \in [\alpha/2, 1 - \alpha/2]$. This is essentially the grid bootstrap of [Hansen, 1999]; however, in our example we do not have to resample the data (to compute the test statistics) for each μ on the grid, but we can sample once and reuse samples in computing the test statistic for each μ. This makes construction of confidence intervals faster. The p-values we get by using the bootstrapped test statistic from (5) are presented in Figure 1.

We can think of computing the bootstrap test using (5) as weighting the bootstrap samples with selection probabilities. Since in more complicated examples computing the selection probabilities efficiently is hard, we devise an alternative way to do bootstrap using wild bootstrap.

Without selection, the wild bootstrap approximates the distribution of $\sqrt{n}(\bar{Y} - \mu)$ with the distribution of $\frac{1}{n}\sum_{i=1}^n (Y_i - \bar{Y})\alpha_i$, where $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_n)$ is a vector of bootstrap weights. We take $\alpha_i \overset{i.i.d.}{\sim} \mathbb{F}_\alpha$ with density f_α. The wild bootstrap approximation of $F_\mu(t)$ becomes

$$\begin{aligned}\widehat{F}^w_\mu(t) &= \mathbb{F}^n_\alpha \times \mathbb{G}\left\{\frac{1}{\sqrt{n}}\sum_{i=1}^n (Y_i - \bar{Y})\alpha_i \leqslant t \,\middle|\, \frac{1}{\sqrt{n}}\sum_{i=1}^n (Y_i - \bar{Y})\alpha_i + \sqrt{n}\mu + \omega > \bar{a}\right\} \\ &= \frac{\mathbb{E}_{\boldsymbol{\alpha}}\left[\bar{G}\left(\bar{a} - \sqrt{n}\mu - \frac{1}{\sqrt{n}}\sum_{i=1}^n (Y_i - \bar{Y})\alpha_i\right)\mathbb{I}_{\left\{\frac{1}{\sqrt{n}}\sum_{i=1}^n (Y_i - \bar{Y})\alpha_i \leqslant t\right\}}\right]}{\mathbb{E}_{\boldsymbol{\alpha}}\left[\bar{G}\left(\bar{a} - \sqrt{n}\mu - \frac{1}{\sqrt{n}}\sum_{i=1}^n (Y_i - \bar{Y})\alpha_i\right)\right]}.\end{aligned} \tag{6}$$

In this example, we can just resample $\boldsymbol{\alpha}$ from $\mathbb{F}^n_\alpha$ or numerically integrate over the numerator and denominator in (6). However, in more complicated selection events in higher dimensions this involves MC techniques.

To compute the wild bootstrap test statistic, we sample weights $\boldsymbol{\alpha} \in \mathbb{R}^n$ from the density proportional to

$$\left(\prod_{i=1}^n f_\alpha(\alpha_i)\right) \cdot \bar{G}\left(\bar{a} - \sqrt{n}\mu - \frac{1}{\sqrt{n}}\sum_{i=1}^n (Y_i - \bar{Y})\alpha_i\right) \tag{7}$$

and compute the test statistic as $\frac{1}{\sqrt{n}}\sum_{i=1}^{n}(Y_i - \bar{Y})\alpha_i$. The pivots are shown in Figure 2, where we use Langevin MC to sample the bootstrap weights from the density in (7). In more complicated examples later the sampling density will also involve constrained optimization variables.

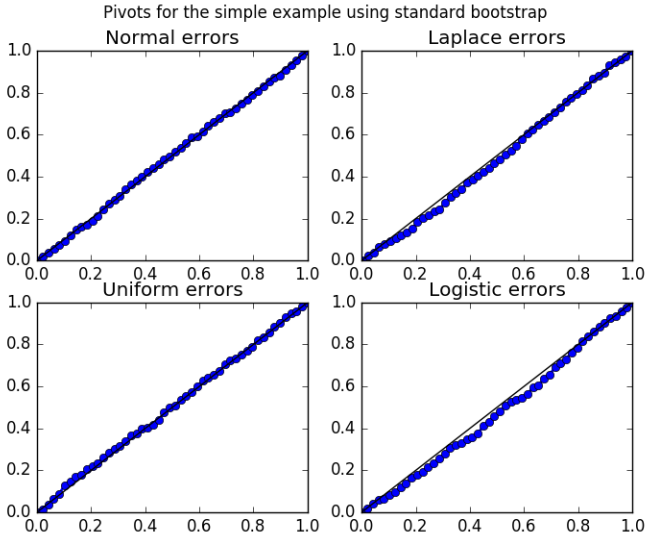

Figure 1: Pivots computed using (5) for $n = 500$ and $B = 5000$ and logistic randomization scaled by 0.8. The distribution of the data is taken to be normal, Laplace, uniform and logistic in four plots shown, each symmetric around its mean $\mu = \frac{1}{\sqrt{n}}$ with variance 1. The test is done for the true mean μ. The threshold is taken to be $\bar{a} = 2$. In all these examples, probability of selection is 11-13%.

2. Selective CLT

In this section, we describe the content of the selective CLT, the fundamental tool used to construct valid tests after selection. Without selection, we are given a statistic $\boldsymbol{T}$ for testing the functional $\boldsymbol{\mu}$ of the unconditional data generating distribution (no selection). In order to construct a test statistic based on $\boldsymbol{\eta}^T\boldsymbol{T}$ to test the functional $\boldsymbol{\eta}^T\boldsymbol{\mu}$ after we select the model and choose $\boldsymbol{\eta}$, we use a pivot, constructed in [Tian and Taylor, 2015b]. Based on their selective CLT result, this pivot is asymptotically Unif$[0, 1]$ under the conditional distribution of the data, conditioning on the selection event. Assuming the logarithm of the randomization density is smooth enough (bounded derivatives up to the third order), we provide the refinements of the selective CLT result using a strictly smaller set of assumptions than [Tian and Taylor, 2015b].

Our selective CLT assumptions are looser in the sense that we do not require the parameter to be close to the selection region (local alternatives) and we prove the Gaussian and nonparametric likelihoods are close.[12] We also generalize their result to allow for testing any vector-valued linear combination of the underlying mean parameter $\boldsymbol{\mu}$. Not only we relax the assumptions of [Tian and Taylor, 2015b], but here we also state the selective CLT for Gaussian randomization under the local alternatives.

Let us now provide the precise setup where we apply the selective CLT. Assume our data at step n is i.i.d. random vectors $\boldsymbol{S}_i, i = 1, \ldots, n$, whose distribution is nonparametric distribution denoted

[12]The latter is stated as an assumption in equation (36), Theorem 9, [Tian and Taylor, 2015b].

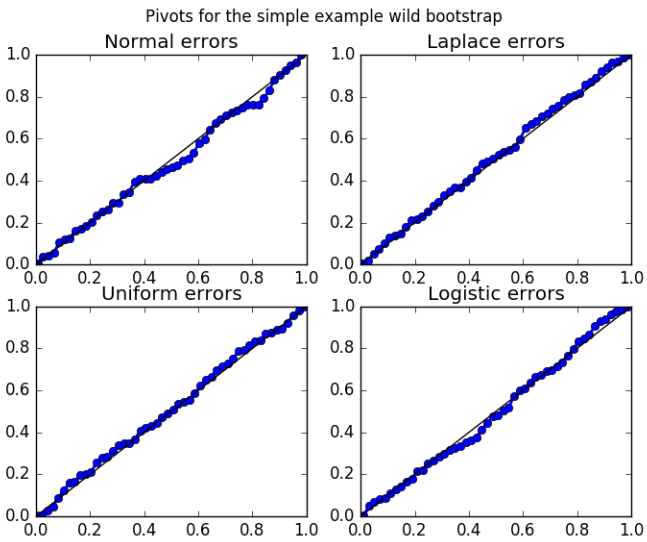

Figure 2: The pivots computed using $\widehat{F}^w_\mu(\sqrt{n}(\bar{Y}-\mu))$ and Langevin sampling from the density in (7) with logistic randomization scaled by 1. The distribution of the data is taken to be normal, Laplace, uniform and logistic in four plots shown, each symmetric around its mean $\mu = \frac{1}{\sqrt{n}}$ with variance 1. The test is done for the hypothesis $\mu = \frac{1}{\sqrt{n}}$. The threshold is taken to be $\bar{a} = 2$. In all these examples, probability of selection is 11-13%.

as $\mathbb{F}_n \in \mathcal{F}_n$. For each n, we take $\mathbb{F}_n \in \mathcal{F}_n$, where $\{\mathcal{F}_n : n \geqslant 1\}$ is a sequence of families of probability distributions (the restrictions on this sequence of families will be made later). Write the data matrix as

$$\boldsymbol{S} = \begin{bmatrix} \boldsymbol{S}_1^T \\ \vdots \\ \boldsymbol{S}_n^T \end{bmatrix}.$$

We are interested in some functional of $\boldsymbol{\mu} = \boldsymbol{\mu}(\mathbb{F}_n) \in \mathbb{R}^p$, where p fixed and does not depend on n throughout. Assume the test statistics $\boldsymbol{T} = \boldsymbol{T}(\boldsymbol{S}) = (T_1(\boldsymbol{S}), \ldots, T_p(\boldsymbol{S})) \in \mathbb{R}^p$, is an asymptotically linear test statistic defined as

$$\boldsymbol{T} = \frac{1}{n}\sum_{i=1}^{n} \boldsymbol{\xi}_i + o_{\mathbb{F}_n}\left(\frac{1}{\sqrt{n}}\right), \tag{AL}$$

where $\boldsymbol{\xi}_i$ are measurable with respect to $\boldsymbol{S}_i$, i.e. $\boldsymbol{\xi}_i = \xi(\boldsymbol{S}_i)$ for some function ξ, and $\mathbb{E}_{\mathbb{F}_n}\boldsymbol{\xi}_i = \boldsymbol{\mu}$ and $\mathrm{Var}(\boldsymbol{\xi}_i) = \boldsymbol{\Sigma}$ for all $i = 1, \ldots, n$.

This implies $\boldsymbol{T}$ satisfies the following uniform CLT pre-selection:

$$\lim_{n\to\infty} \sup_{\mathbb{F}_n \in \mathcal{F}_n} \sup_{\boldsymbol{t}\in\mathbb{R}^p} \left| \mathbb{F}_n\left\{\sqrt{n}\left(\boldsymbol{T} - \boldsymbol{\mu}(\mathbb{F}_n)\right) \leqslant \boldsymbol{t}\right\} - \mathbb{P}_{\boldsymbol{G}\sim\mathcal{N}_p(\boldsymbol{0},\boldsymbol{\Sigma})}\{\boldsymbol{G} \leqslant \boldsymbol{t}\}\right| = 0. \tag{CLT}$$

This is a CLT statement uniform across $\boldsymbol{t} \in \mathbb{R}^p$ and uniform across the class of distributions $\mathbb{F}_n \in \mathcal{F}_n$. Since uniform convergence statements lead to uniformly consistent confidence intervals, we will also prove the uniform bootstrap convergence statements. Hence it holds for any sequence of distributions $\{\mathbb{F}_n \in \mathcal{F}_n : n \geqslant 1\}$.

Remark 3 *The selective CLT of [Tian and Taylor, 2015b] assumes the stronger condition above holds, i.e.* (AL) *instead of just* (CLT)*. By using the method of [Chatterjee, 2005] they also get*

the rate of convergence and not just asymptotic convergence. Furthermore, they do not assume the sequence of parameters converges as the the sample size n increases (as the assumption in [Tibshirani et al., 2015] although the latter does non-randomized case) but the local alternatives condition (see below).

Assuming the sequence of parameters converges, we prove a version of selective CLT using the condition (CLT)*, convergence of the sequence of selection regions any continuity assumption on the randomization distribution in Appendix A. However, since this result does not allow for rare events, we below show a selective CLT version allowing for underlying parameters to be far from the selection region.*

Remark 4 *Having uniformly consistent confidence intervals and not just consistent is an important inferential goal since it makes a difference in coverage guarantees for finite n [Romano et al., 2012], [Tibshirani et al., 2015].*

Remark 5 *There are several frameworks in selective inference literature: saturated ([Lee et al., 2016], [Loftus and Taylor, 2015], [Tibshirani et al., 2015]) and selected model framework ([Fithian et al., 2014]). In the saturated model setting, the families $\{\mathcal{F}_n : n \geqslant 1\}$ are fixed in advance, i.e. chosen without looking at the data. There is no model being selected here, but the parameter of interest, $\boldsymbol{\eta}^T\boldsymbol{\mu}(\mathbb{F}_n)$ ($\boldsymbol{\eta} \in \mathbb{R}^p$), depends on the selected model. Hence, conditioning on the selection event accounts for the choice of the functional for which to report a confidence interval. In the selected model setting, the sequence of families $\{\mathcal{F}_n : n \geqslant 1\}$ could have been generated based on some model selection procedure. In this case, the choice of families is data dependent. The conditional event accounts for the choice of both the family, $\mathcal{F}_n$, and the functional, $\boldsymbol{\eta}^T\boldsymbol{\mu}$. As in [Tian and Taylor, 2015b], we allow both saturated and selected model setting, however, we focus here on the saturated model setting where the inference is done for the population parameters chosen by looking at the selected model.*

2.1. Linear decomposition

Since we are interested in testing several linear combinations of $\boldsymbol{\mu}$ at once, let our vector-valued parameter of interest be $\boldsymbol{A}^T\boldsymbol{\mu}$ for some matrix $\boldsymbol{A} \in \mathbb{R}^{p\times a}$. Denote also

$$\boldsymbol{\Sigma}_{\boldsymbol{A}} = \boldsymbol{A}^T\boldsymbol{\Sigma}\boldsymbol{A}, \quad \boldsymbol{C} = \boldsymbol{\Sigma}\boldsymbol{A}\boldsymbol{\Sigma}_{\boldsymbol{A}}^{-1}, \quad \boldsymbol{T}_{\boldsymbol{A}} = \left(\boldsymbol{I}_p - \boldsymbol{C}\boldsymbol{A}^T\right)\boldsymbol{T},$$

where $\boldsymbol{I}_p$ is the identity matrix of dimension p. We have the decomposition

$$\boldsymbol{T} = \boldsymbol{T}_{\boldsymbol{A}} + \boldsymbol{C}\boldsymbol{A}^T\boldsymbol{T}. \tag{8}$$

If the vector $\boldsymbol{T}$ were normally distributed $\mathcal{N}_p\left(\boldsymbol{\mu}, \frac{1}{n}\boldsymbol{\Sigma}\right)$ with known $\boldsymbol{\Sigma}$ and unknown $\boldsymbol{\mu}$ and if $\boldsymbol{A}^T\boldsymbol{\mu}$ was the parameter of interest the sufficient statistic for the nuisance parameters $\mathcal{P}_{\boldsymbol{A}}^{\perp}\boldsymbol{\mu}$ would be $\boldsymbol{T}_{\boldsymbol{A}}$. In this case, we would base the inference on the conditional law of $\boldsymbol{A}^T\boldsymbol{T}$ given $\boldsymbol{T}_{\boldsymbol{A}}$ and the selection region. This construction of the pivot in exponential families is a standard construction which was used in [Fithian et al., 2014] for selective inference. Under a CLT, the limiting family pre-selection is a Gaussian exponential family, hence its selective counterpart is a Gaussian exponential family subject to selection. This family is used to construct the pivotal quantity. Indeed, the main feature of selective inference is that under many interesting scenarios the limiting model is *not* Gaussian but a Gaussian model subject to selection. This observation suggests natural test statistics derived under the limiting Gaussian selective model. Much work is often known about the Gaussian model pre-selection, and this generally transfers to the selective model.

The idea of how to construct the pivot is taken from the exponential families and later proved to work in non-parametric settings as well.

We now assume the model selection event is based on $\boldsymbol{T}$ and satisfies the affine constrains. Hence

the selection region at step n is denoted as

$$\begin{aligned}\text{selection}(\boldsymbol{T},\boldsymbol{\omega}) &:= \left\{\sqrt{n}\boldsymbol{A}_M\boldsymbol{T}+\boldsymbol{\omega}\in\boldsymbol{H}_M\right\}\\ &= \left\{\boldsymbol{A}_M\left(\sqrt{n}\left(\boldsymbol{T}-\boldsymbol{\mu}\right)\right)+\boldsymbol{A}_M\sqrt{n}\boldsymbol{\mu}+\boldsymbol{\omega}\in\boldsymbol{H}_M\right\}\\ &= \left\{\sqrt{n}\boldsymbol{A}_M\boldsymbol{T}_{\boldsymbol{A}}+\boldsymbol{A}_M\boldsymbol{C}\left(\sqrt{n}\left(\boldsymbol{A}^T\boldsymbol{T}-\boldsymbol{A}^T\boldsymbol{\mu}\right)\right)+\sqrt{n}\boldsymbol{A}_M\boldsymbol{C}\boldsymbol{A}^T\boldsymbol{\mu}+\boldsymbol{\omega}\in\boldsymbol{H}_M\right\},\end{aligned}$$

where $\boldsymbol{A}_M\in\mathbb{R}^{d\times p}$, $\boldsymbol{\omega}\in\mathbb{R}^d$ is a randomization noise and $\boldsymbol{H}_M\subset\mathbb{R}^d$. Assume $\boldsymbol{\omega}\sim\mathbb{G}$, where $\mathbb{G}$ is a known distribution in $\mathbb{R}^d$ specified by the user. Further assume the noise $\boldsymbol{\omega}$ is independent of the data $\boldsymbol{S}$. Also, d is assumed to be fixed and does not depend on n. Denoting $\boldsymbol{\Delta}=\boldsymbol{\Delta}(\mathbb{F}_n)=\sqrt{n}\boldsymbol{\mu}$ and $\boldsymbol{Z}=\sqrt{n}\left(\boldsymbol{T}-\boldsymbol{\mu}\right)$, we write the selection region also as

$$\begin{aligned}\text{selection}(\boldsymbol{T},\boldsymbol{\omega}) &= \{\boldsymbol{A}_M(\boldsymbol{Z}+\boldsymbol{\Delta})+\boldsymbol{\omega}\in\boldsymbol{H}_M\}\\ &= \left\{\boldsymbol{A}_M\boldsymbol{Z}_{\boldsymbol{A}}+\boldsymbol{A}_M\boldsymbol{C}\boldsymbol{A}^T\boldsymbol{Z}+\boldsymbol{A}_M\boldsymbol{\Delta}+\boldsymbol{\omega}\in\boldsymbol{H}_M\right\},\end{aligned}$$

where $\boldsymbol{Z}_{\boldsymbol{A}}=\left(\boldsymbol{I}_p-\boldsymbol{C}\boldsymbol{A}^T\right)\boldsymbol{Z}$.

We assume our data is generated from the distribution $\mathbb{F}_n^*$, which is the marginal distribution of the data $\boldsymbol{S}\sim\mathbb{F}_n$ conditional on the randomized selection event $\text{selection}(\boldsymbol{T},\boldsymbol{\omega})$, $\boldsymbol{\omega}\sim\mathbb{G}$. All the guarantees given will be under this distribution.

To introduce the pivot from [Tian and Taylor, 2015b], let us first define a survival function

$$\begin{aligned}&\mathcal{P}\left(t,\boldsymbol{G}_{\boldsymbol{A}},\boldsymbol{A}^T\boldsymbol{\mu}\right) := \Phi_n\times\mathbb{G}\left\{\|\boldsymbol{\Sigma}_{\boldsymbol{A}}^{-1/2}\boldsymbol{A}^T(\sqrt{n}(\boldsymbol{G}-\boldsymbol{\mu}))\|_2\geqslant t\,\middle|\,\text{selection}(\boldsymbol{G},\boldsymbol{\omega}),\boldsymbol{G}_{\boldsymbol{A}}\right\}\\ &=\frac{\int_{\|\boldsymbol{\Sigma}_{\boldsymbol{A}}^{-1/2}\boldsymbol{A}^T\boldsymbol{\tau}\|_2\geqslant t}\mathbb{G}\{\boldsymbol{\omega}:\sqrt{n}\boldsymbol{A}_M\boldsymbol{G}_{\boldsymbol{A}}+\boldsymbol{A}_M\boldsymbol{C}\boldsymbol{\tau}+\sqrt{n}\boldsymbol{A}_M\boldsymbol{C}\boldsymbol{A}^T\boldsymbol{\mu}+\boldsymbol{\omega}\in\boldsymbol{H}_M\}\phi_a(\boldsymbol{\tau};\boldsymbol{0},\boldsymbol{\Sigma}_{\boldsymbol{A}})d\boldsymbol{\tau}}{\int_{\mathbb{R}^a}\mathbb{G}\{\boldsymbol{\omega}:\sqrt{n}\boldsymbol{A}_M\boldsymbol{G}_{\boldsymbol{A}}+\boldsymbol{A}_M\boldsymbol{C}\boldsymbol{\tau}+\sqrt{n}\boldsymbol{A}_M\boldsymbol{C}\boldsymbol{A}^T\boldsymbol{\mu}+\boldsymbol{\omega}\in\boldsymbol{H}_M\}\phi_a(\boldsymbol{\tau};\boldsymbol{0};\boldsymbol{\Sigma}_{\boldsymbol{A}})d\boldsymbol{\tau}},\end{aligned}$$

where the probability on the RHS is under $\boldsymbol{G}\sim\mathcal{N}_p\left(\boldsymbol{\mu},\frac{\boldsymbol{\Sigma}}{n}\right):=\Phi_n$ and under $\boldsymbol{\omega}\sim\mathbb{G}$. Here, $\phi_a(\cdot;\boldsymbol{0},\boldsymbol{\Sigma}_{\boldsymbol{A}})$ is the density of $\mathcal{N}_a(\boldsymbol{0},\boldsymbol{\Sigma}_{\boldsymbol{A}})$. Note that the conditioning on the RHS is on the selection event for $(\boldsymbol{G},\boldsymbol{\omega})$ and the value of $\boldsymbol{G}_{\boldsymbol{A}}:=\left(\boldsymbol{I}_p-\boldsymbol{C}\boldsymbol{A}^T\right)\boldsymbol{G}$, which is the sufficient statistic for the nuisance parameters assuming our parameter of interest is $\boldsymbol{A}^T\boldsymbol{\mu}$. Note that computing this pivot requires knowing $\boldsymbol{\Sigma}$. We will assume the variance $\boldsymbol{\Sigma}$ is known. However, having uniformly consistent estimate of variance would suffice.[13]

Denoting the conditional distribution of $\boldsymbol{G}$ as Φ_n^*, which is the distribution of $\boldsymbol{G}\sim\Phi_n$ conditional on the event $\text{selection}(\boldsymbol{G},\boldsymbol{\omega})$, it is not hard to see that $\mathcal{P}\left(\sqrt{n}\boldsymbol{A}^T(\boldsymbol{G}-\boldsymbol{\mu}),\boldsymbol{G}_{\boldsymbol{A}},\boldsymbol{A}^T\boldsymbol{\mu}\right)$ is distributed as Unif$[0,1]$ under this conditional distribution of $\boldsymbol{G}$ (Lemma 7 in [Tian and Taylor, 2015b]). Hence, under normality assumptions, the constructed pivot is exact and in what follows we state the relaxed result.

In order to show that that the pivot defined as

$$\mathcal{P}\left(\sqrt{n}(\boldsymbol{A}^T\boldsymbol{T}-\boldsymbol{A}^T\boldsymbol{\mu}),\boldsymbol{T}_{\boldsymbol{A}},\boldsymbol{A}^T\boldsymbol{\mu}\right).$$

for non-normal $\boldsymbol{T}$ is asymptotically Unif$[0,1]$, [Tian and Taylor, 2015b] make the following assumptions. Note also that [Tian and Taylor, 2015b] consider only 1-dimensional parameters, i.e $A\in\mathbb{R}^{p\times 1}$.

- Uniformly bounded MGF in some neighborhood of zero: Precisely, assume there exists $\tau>0$ and a constant C_τ such that

$$\sup_{n\geqslant 1}\sup_{\mathbb{F}_n\in\mathcal{F}_n}\mathbb{E}_{\boldsymbol{s}\sim\mathbb{F}_n}\left[e^{\tau\|\boldsymbol{\xi}(\boldsymbol{s})-\boldsymbol{\mu}\|_1}\right]\leqslant C_\tau. \tag{MGF}$$

[13] Lemma 14 in [Tian and Taylor, 2015b] states that under some integrability conditions the pivot can be computed with the estimated variance.

- Gaussian (Φ_n) and non-parametric ($\mathbb{F}_n$) selective likelihood ratios are asymptotically close:

$$\frac{|(\mathbb{F}_n \times \mathbb{G})\{\text{selection}(\boldsymbol{T}, \boldsymbol{\omega})\} - (\Phi_n \times \mathbb{G})\{\text{selection}(\boldsymbol{G}, \boldsymbol{\omega})\}|}{(\Phi_n \times \mathbb{G})\{\text{selection}(\boldsymbol{G}, \boldsymbol{\omega})\}} = O\left(\frac{1}{\sqrt{n}}\right) \tag{LR}$$

  where $\boldsymbol{G} \sim \mathcal{N}_p\left(\boldsymbol{\mu}, \frac{1}{n}\boldsymbol{\Sigma}\right) = \Phi_n$. This condition ensures

$$\lim_{n\to\infty} \mathbb{E}_{\mathbb{F}_n}\left[|\ell_{\mathbb{F}_n}(\boldsymbol{T}) - \ell_{\Phi_n}(\boldsymbol{T})|\right] = 0,$$

  where

$$\ell_{\mathbb{F}_n}(\boldsymbol{T}) = \frac{d\mathbb{F}_n^*}{d\mathbb{F}_n}(\boldsymbol{T}) := \frac{\mathbb{G}\{\text{selection}(\boldsymbol{T}, \boldsymbol{\omega})\}}{(\mathbb{F}_n \times \mathbb{G})\{\text{selection}(\boldsymbol{T}', \boldsymbol{\omega})\}}$$

  and

$$\ell_{\Phi_n}(\boldsymbol{T}) = \frac{d\Phi_n^*}{d\Phi_n}(\boldsymbol{T}) := \frac{\mathbb{G}\{\text{selection}(\boldsymbol{T}, \boldsymbol{\omega})\}}{(\Phi_n \times \mathbb{G})\{\text{selection}(\boldsymbol{T}', \boldsymbol{\omega})\}}$$

  are the selective likelihood ratios when the data is distributed from $\mathbb{F}_n$ and Φ_n, respectively.
- Local alternatives: Assume there is some norm $\|\cdot\|_h$ in $\mathbb{R}^d$ and the corresponding distance $d_h(\cdot,\cdot)$ induced by this norm such that

$$d_h\left(0, \boldsymbol{H}_M - \sqrt{n}\boldsymbol{A}_M\boldsymbol{\mu}(\mathbb{F}_n)\right) \leqslant B, \tag{LA}$$

  from some constant B. This condition does not allow for the underlying parameter to be far from the selection region, hence the selection events cannot have low probability (cannot be rare).
- Lower bound on the randomization: There exists a constant C^- such that

$$\mathbb{G}\left\{\boldsymbol{\omega} \in \boldsymbol{H}_M - \boldsymbol{\theta}\right\} \geqslant C^- e^{-d_h(\mathbf{0}, \boldsymbol{H}_M - \boldsymbol{\theta})}, \quad \forall \boldsymbol{\theta} \in \mathbb{R}^d. \tag{LB}$$

  An example of randomization satisfying this condition is Logistic distribution. Note that the Gaussian randomization does not satisfy this assumption.
- Smoothness of the randomization: Suppose $\mathbb{G}$ has a density g with integrable first three derivatives,

$$\int_{\mathbb{R}^d} \left\|\partial^j g(\boldsymbol{\omega})\right\|_\infty d\boldsymbol{\omega} \leqslant C_j, \quad j = 0, 1, 2, 3. \tag{SR}$$

- The norm of matrix $\boldsymbol{A}_M$ does not grow with n, i.e. there exists a finite constant C_M such that

$$\limsup_{n\to\infty} \|\boldsymbol{A}_M\|_{h,2} \leqslant C_M < \infty, \tag{NA}$$

  where the matrix norm $\|\cdot\|_{h,2}$ is defined as $\|\boldsymbol{A}_M\|_{h,2} = \sup_{\boldsymbol{u}\in\mathbb{R}^p} \frac{\|\boldsymbol{A}_M\boldsymbol{u}\|_h}{\|\boldsymbol{u}\|_2}$.

Under all the assumptions above (including (AL)), [Tian and Taylor, 2015b] show the pivot $\mathcal{P}$ (constructed for estimating a linear functional $\boldsymbol{\eta}^T\boldsymbol{\mu}$) is asymptotically Unif$[0,1]$ in the non-parametric setting under $\boldsymbol{S} \sim \mathbb{F}_n^*$. They also provide $O(1/\sqrt{n})$ rate of convergence.

2.2. Selective CLT under Lipschitz randomization

We show a selective CLT result with the same guarantees assuming the randomization $\boldsymbol{\omega} \in \mathbb{R}^d$ comes from a distribution $\mathbb{G}$ on $\mathbb{R}^d$ with a density g

$$g(\boldsymbol{\omega}) = \frac{1}{C_g}\exp(-\tilde{g}(\boldsymbol{\omega})),$$

where we assume $\tilde{g}$ has bounded derivatives up to order at least 3 (call the bound on all these derivatives K_g).

More precisely, for any multi-index $\boldsymbol{\alpha} \in \mathbb{N}_0^d$ with $\|\boldsymbol{\alpha}\|_1 \leqslant 3$ we have

$$|\partial_{\boldsymbol{\omega}}^{\boldsymbol{\alpha}} \tilde{g}(\boldsymbol{\omega})| \leqslant K_g \tag{S}$$

for some constant K_g. This implies the Lipschitz property of $\tilde{g}$,

$$|\tilde{g}(\boldsymbol{x}) - \tilde{g}(\boldsymbol{y})| \leqslant K_g \|\boldsymbol{x} - \boldsymbol{y}\|_h, \ \forall \boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^d, \tag{Lip}$$

where $\|\cdot\|_h$ is some norm in $\mathbb{R}^d$. Note that in particular that the logarithm of Gaussian density fails to satisfy the above assumptions as its gradient is unbounded over $\mathbb{R}^d$. Hence, one of the main properties we require is the the logarithm of the density is globally Lipschitz. Bounded second and third derivatives are required in our proof though these assumptions could likely be relaxed.

Using the assumption (S) on the randomization, we prove (LR) holds (Lemma 17) and the selective CLT (Theorem 18) both for finite n. Here we state the asymptotic version of Theorem 18 as a corollary. All the proofs for this section are done in detail in Appendix B.

Corollary 6 (Lipschitz randomization: Selective CLT) *Assuming* (AL), (MGF), (NA) *and* (S), *we have*

$$\lim_{n\to\infty} \sup_{\mathbb{F}_n \in \mathcal{F}_n} \sup_{t\in[0,1]} \left| \mathbb{F}_n^* \left\{ \mathcal{P}\left(\sqrt{n}(\boldsymbol{A}^T\boldsymbol{T} - \boldsymbol{A}^T\boldsymbol{\mu}), \boldsymbol{T}_{\boldsymbol{A}}, \boldsymbol{A}^T\boldsymbol{\mu}\right)\right\} - t \right| = 0. \tag{P}$$

2.3. Selective CLT with Gaussian randomization

In this section, we state the selective CLT result for the Gaussian randomization. The Gaussian randomization arises naturally in the data carving framework (see Section 7), where we do not take randomization independent of the data from a known distribution (as until now) but we compute it from the data and the random split indicators. Assuming slightly weaker condition than the local alternatives, that the parameter value is at the distance at most growing like $o(\log n)$ from the selection region, and sub-Gaussian data we obtain a selective CLT result for the Gaussian randomization that we now describe in more detail.

We take the randomization distribution to be Gaussian in d dimensions

$$\mathbb{G} = \mathcal{N}_d(\mathbf{0}, c\boldsymbol{I}_d). \tag{G}$$

Assuming the matrix $\boldsymbol{A}$ is of size $d \times 1$ and the variance of $\boldsymbol{A}^T\boldsymbol{Z} =: T$ is 1, the pivot becomes

$$\frac{\int_{\boldsymbol{H}_M} \bar{\Phi}\left(\sqrt{\frac{\|\boldsymbol{v}\|_2^2+c}{c}}T - \frac{\boldsymbol{v}^T\boldsymbol{w}}{\sqrt{c(\|\boldsymbol{v}\|_2^2+c)}}\right) \cdot \exp\left(-\frac{1}{2}\boldsymbol{w}^T(\boldsymbol{v}\boldsymbol{v}^T + c\boldsymbol{I})^{-1}\boldsymbol{w}\right) d\boldsymbol{\omega}'}{\int_{\boldsymbol{H}_M} \exp\left(-\frac{1}{2}\boldsymbol{w}^T(\boldsymbol{v}\boldsymbol{v}^T + c\boldsymbol{I})^{-1}\boldsymbol{w}\right) d\boldsymbol{\omega}'},$$

where $\boldsymbol{v} = \boldsymbol{A}_M\boldsymbol{C}$ and $\boldsymbol{w} = \boldsymbol{\omega}' - \boldsymbol{A}_M\boldsymbol{z}_{\boldsymbol{A}} - \boldsymbol{A}_M\boldsymbol{\Delta}$ (derived in the Appendix C.1).

In order to prove the above test statistic is uniform $[0,1]$ asymptotically, we make the following additional assumptions here.

- The selection region is rectangular $\boldsymbol{H}_M = \prod_{i=1}^n [b_i, \infty)$ for some vector $\boldsymbol{b} \in \mathbb{R}^d$.
- The distance of the parameter to the selection region is not growing too fast with the sample size:

$$d(\mathbf{0}, \boldsymbol{H}_M - \boldsymbol{A}_M\boldsymbol{\Delta}) = o(\log n). \tag{LA-w}$$

- The random variables coming from $\mathbb{F}_n$ are sub-Gaussian, i.e. there exists $b > 0$ such that

$$\sup_{n\geqslant 1} \sup_{\mathbb{F}_n \in \mathcal{F}_n} \mathbb{E}_{\boldsymbol{s}\sim\mathbb{F}_n}\left[e^{b\|\boldsymbol{\xi}(\boldsymbol{s})-\boldsymbol{\mu}\|_2^2}\right] \leqslant C_b. \tag{subG}$$

We are ready to state the Selective CLT for the Gaussian randomization, which is a corollary of the Theorem 23 proved in Appendix C. The result relies on two important lemmas, Lemma 19 and Lemma 20 in the Appendix C, establishing the 'smoothness' of the likelihood and the pivot.

Corollary 7 (Gaussian randomization: Selective CLT) *Assuming* (AL), (LA-w), (G), (NA) *and* (subG), *we have*

$$\lim_{n\to\infty} \sup_{\mathbb{F}_n\in\mathcal{F}_n} \sup_{t\in[0,1]} \left|\mathbb{F}_n^*\left\{\mathcal{P}\left(\sqrt{n}(\boldsymbol{A}^T\boldsymbol{T}-\boldsymbol{A}^T\boldsymbol{\mu}),\boldsymbol{T}_{\boldsymbol{A}},\boldsymbol{A}^T\boldsymbol{\mu}\right)\right\}-t\right|=0. \tag{9}$$

3. Consistency after selection

This main result of this section is to prove that the estimators consistent for zero under the original measure without selection, $\mathbb{F}_n$, are also consistent for zero under the $\mathbb{F}_n^*$, the marginal distribution of the data conditional on the selection event. We do that by analyzing the selective likelihood ratio of these two distributions.

Using the introduced notation, the selective likelihood ratio becomes

$$\ell_{\mathbb{F}_n}(\boldsymbol{t}) := \frac{d\mathbb{F}_n^*(\boldsymbol{t})}{d\mathbb{F}_n(\boldsymbol{t})}, \quad \boldsymbol{t}\in\mathbb{R}^p.$$

which in our problem becomes

$$\begin{aligned}
\ell_{\mathbb{F}_n}(\boldsymbol{t}) &= \frac{\mathbb{G}\left\{\boldsymbol{\omega}\in\mathbb{R}^d : (\boldsymbol{t},\boldsymbol{\omega}) \text{ selected}\right\}}{(\mathbb{F}_n\times\mathbb{G})\left\{(\boldsymbol{Z},\boldsymbol{\omega})\in\mathbb{R}^p\times\mathbb{R}^d : (\boldsymbol{Z},\boldsymbol{\omega}) \text{ selected}\right\}} \\
&= \frac{\mathbb{G}\left\{\boldsymbol{\omega} : \boldsymbol{A}_M\boldsymbol{z}+\boldsymbol{A}_M\boldsymbol{\Delta}+\boldsymbol{\omega}\in\boldsymbol{H}_M\right\}}{(\mathbb{F}_n\times\mathbb{G})\left\{(\boldsymbol{Z},\boldsymbol{\omega}) : \boldsymbol{A}_M\boldsymbol{Z}+\boldsymbol{A}_M\boldsymbol{\Delta}+\boldsymbol{\omega}\in\boldsymbol{H}_M\right\}} \\
&= \frac{\mathbb{G}\left\{\boldsymbol{H}_M-\boldsymbol{A}_M\boldsymbol{z}-\boldsymbol{A}_M\boldsymbol{\Delta}\right\}}{\mathbb{E}_{\boldsymbol{Z}\sim\mathbb{F}_n}\left[\mathbb{G}\left\{\boldsymbol{H}_M-\boldsymbol{A}_M\boldsymbol{\Delta}-\boldsymbol{A}_M\boldsymbol{Z}\mid\boldsymbol{Z}\right\}\right]}
\end{aligned}$$

($\boldsymbol{z}=\sqrt{n}(\boldsymbol{t}-\boldsymbol{\mu})$). Denoting

$$Q_1(\boldsymbol{z};\boldsymbol{\Delta}) = \mathbb{G}\left\{\boldsymbol{H}_M-\boldsymbol{A}_M\boldsymbol{\Delta}-\boldsymbol{A}_M\boldsymbol{z}\right\}, \quad \boldsymbol{z}\in\mathbb{R}^p,$$

we write the likelihood ratio as

$$\ell_{\mathbb{F}_n}(\boldsymbol{t}) = \frac{Q_1(\boldsymbol{z};\boldsymbol{\Delta})}{\mathbb{E}_{\mathbb{F}_n}\left[Q_1(\boldsymbol{Z};\boldsymbol{\Delta})\right]}.$$

The following lemma is providing an integrable upper bounded on the selective likelihood ratio and is used later to show the consistency result.

Lemma 8 (Upper bound on the selective likelihood ratio) *Using the* (Lip) *assumption, we have*

$$\ell_{\mathbb{F}_n}(\boldsymbol{t}) \leqslant \frac{e^{K_g\|\boldsymbol{A}_M\boldsymbol{z}\|_h}}{\mathbb{E}_{\mathbb{F}_n}\left[e^{-K_g\|\boldsymbol{A}_M\boldsymbol{Z}\|_h}\right]}. \tag{10}$$

Proof Using the change of variables $\boldsymbol{\omega}'=\boldsymbol{\omega}+\boldsymbol{A}_M\boldsymbol{z}+\boldsymbol{A}_M\boldsymbol{\Delta}$ and the Lipschitz assumption on $\tilde{g}$,

we have the lower bound

$$\begin{aligned}
Q_1(\boldsymbol{z};\boldsymbol{\Delta}) &= \int\limits_{\boldsymbol{\omega}\in\boldsymbol{H}_M-\boldsymbol{A}_M\boldsymbol{z}-\boldsymbol{A}_M\boldsymbol{\Delta}} \frac{1}{C_g}\exp(-\tilde{g}(\boldsymbol{\omega}))d\boldsymbol{\omega} \\
&= \frac{1}{C_g}\int\limits_{\boldsymbol{\omega}'\in\boldsymbol{H}_M} \exp\left(-\tilde{g}(\boldsymbol{\omega}'-\boldsymbol{A}_M\boldsymbol{z}-\boldsymbol{A}_M\boldsymbol{\Delta})\right)d\boldsymbol{\omega}' \\
&\geqslant \frac{1}{C_g}\int\limits_{\boldsymbol{\omega}'\in\boldsymbol{H}_M} \exp\left(-\tilde{g}(\boldsymbol{\omega}'-\boldsymbol{A}_M\boldsymbol{\Delta})-K_g\|\boldsymbol{A}_M\boldsymbol{z}\|_h\right)d\boldsymbol{\omega}' \\
&= e^{-K_g\|\boldsymbol{A}_M\boldsymbol{z}\|_h}\int\limits_{\boldsymbol{\omega}'\in\boldsymbol{H}_M} \frac{1}{C_g}\exp\left(-\tilde{g}(\boldsymbol{\omega}'-\boldsymbol{A}_M\boldsymbol{\Delta})\right)d\boldsymbol{\omega}' \\
&= Q_1(\boldsymbol{0};\boldsymbol{\Delta})\cdot e^{-K_g\|\boldsymbol{A}_M\boldsymbol{z}\|_h},
\end{aligned}$$

implying

$$\mathbb{E}_{\mathbb{F}_n}\left[Q_1(\boldsymbol{Z};\boldsymbol{\Delta})\right] \geqslant Q_1(\boldsymbol{0};\boldsymbol{\Delta})\cdot\mathbb{E}_{\mathbb{F}_n}\left[e^{-K_g\|\boldsymbol{A}_M\boldsymbol{Z}\|_h}\right].$$

Similarly, we have an upper bound

$$\begin{aligned}
Q_1(\boldsymbol{z};\boldsymbol{\Delta}) &\leqslant \frac{1}{C_g}\int\limits_{\boldsymbol{\omega}'\in\boldsymbol{H}_M} \exp\left(-\tilde{g}(\boldsymbol{\omega}'-\boldsymbol{A}_M\boldsymbol{\Delta})+K_g\|\boldsymbol{A}_M\boldsymbol{z}\|_h\right)d\boldsymbol{\omega}' \\
&= e^{K_g\|\boldsymbol{A}_M\boldsymbol{z}\|_h}\int\limits_{\boldsymbol{\omega}'\in\boldsymbol{H}_M} \frac{1}{C_g}\exp\left(-\tilde{g}(\boldsymbol{\omega}'-\boldsymbol{A}_M\boldsymbol{\Delta})\right)d\boldsymbol{\omega}' \\
&= Q_1(\boldsymbol{0};\boldsymbol{\Delta})\cdot e^{K_g\|\boldsymbol{A}_M\boldsymbol{z}\|_h}.
\end{aligned}$$

■

The following lemma shows that if a sequence of estimators is consistent for zero under the original distribution of the data ($\boldsymbol{S}\sim\mathbb{F}_n$) then that sequence is also consistent for zero under the conditional distribution ($\boldsymbol{S}\sim\mathbb{F}_n^*$). The proof is similar to the proof of Lemma 1 in [Tian and Taylor, 2015b], hence the assumptions of the Lemma may be thought of as a way to verify the integrability assumptions of that result. The lemma is used later in proving the consistency of bootstrap statistic under the conditional distribution.

Lemma 9 (Consistency) *Assume* (CLT), (Lip) *and* (NA) *hold. Let* $f_n(\boldsymbol{T})$ *be a sequence of estimators such that*

$$f_n(\boldsymbol{T})\overset{\mathbb{F}_n}{\to}0$$

as $n\to\infty$ *uniformly across* $\mathbb{F}_n\in\mathcal{F}_n$, *i.e. for every* $\epsilon>0$, $\lim\limits_{n\to\infty}\sup\limits_{\mathbb{F}_n\in\mathcal{F}_n}\mathbb{F}_n\left\{f_n(\boldsymbol{T})>\epsilon\right\}=0$. *Then we also have*

$$f_n(\boldsymbol{T})\overset{\mathbb{F}_n^*}{\to}0$$

as $n\to\infty$ *uniformly across* $\mathbb{F}_n\in\mathcal{F}_n$, *i.e. for every* $\epsilon>0$, $\lim\limits_{n\to\infty}\sup\limits_{\mathbb{F}_n\in\mathcal{F}_n}\mathbb{F}_n^*\left\{f_n(\boldsymbol{T})>\epsilon\right\}=0$.

Proof

Take any $\epsilon > 0$. Using the upper bound on the selective likelihood ratio bound from (10), we get the following inequality

$$\begin{aligned}\mathbb{F}_n^*\{f_n(\boldsymbol{T}) > \epsilon\} &= \mathbb{F}_n\left\{f_n(\boldsymbol{T}) > \epsilon \,\middle|\, \text{selection}(\boldsymbol{T}, \boldsymbol{\omega})\right\} \\ &= \int \mathbb{I}_{\{f_n(\boldsymbol{t})>\epsilon\}} d\mathbb{F}_n^*(\boldsymbol{t}) = \int \mathbb{I}_{\{f_n(\boldsymbol{t})>\epsilon\}} \ell_{\mathbb{F}_n}(\boldsymbol{t}) d\mathbb{F}_n(\boldsymbol{t}) \\ &\leqslant \frac{1}{\mathbb{E}_{\mathbb{F}_n}\left[e^{-K_g\|\boldsymbol{A}_M\boldsymbol{Z}\|_h}\right]} \int \mathbb{I}_{\{f_n(\boldsymbol{t})>\epsilon\}} e^{K_g\|\boldsymbol{A}_M\boldsymbol{z}\|_h} d\mathbb{F}_n(\boldsymbol{t}).\end{aligned}$$

Using (CLT) and (NA) assumption we have $\|\boldsymbol{A}_M\boldsymbol{z}\|_h = O_{\mathbb{F}_n}(1)$, hence for every $\delta > 0$, there exists C_δ such that

$$\mathbb{F}_n\{\|\boldsymbol{A}_M\boldsymbol{Z}\|_g > C_\delta\} < \delta$$

for n sufficiently large. Since

$$\begin{aligned}&\int \mathbb{I}_{\{f_n(\boldsymbol{t})>\epsilon\}} e^{K_g\|\boldsymbol{A}_M\boldsymbol{z}\|_h} d\mathbb{F}_n(\boldsymbol{t}) \\ &= \int \mathbb{I}_{\{f_n(\boldsymbol{t})>\epsilon\}} e^{K_g\|\boldsymbol{A}_M\boldsymbol{z}\|_h} \mathbb{I}_{\{\|\boldsymbol{A}_M\boldsymbol{z}\|_h \leqslant C(\delta)\}} d\mathbb{F}_n(\boldsymbol{t}) + \int \mathbb{I}_{\{f_n(\boldsymbol{t})>\epsilon\}} e^{K_g\|\boldsymbol{A}_M\boldsymbol{z}\|_h} \mathbb{I}_{\{\|\boldsymbol{A}_M\boldsymbol{z}\|_h > C_\delta\}} d\mathbb{F}_n(\boldsymbol{t}) \\ &\leqslant e^{K_g\cdot C_\delta}\mathbb{F}_n\{f_n(\boldsymbol{T}) > \epsilon\} + \mathbb{E}_{\mathbb{F}_n}\left[e^{K_g\|\boldsymbol{A}_M\boldsymbol{Z}\|_h}\mathbb{I}_{\{\|\boldsymbol{A}_M\boldsymbol{Z}\|_h > C_\delta\}}\right] \\ &\leqslant e^{K_g\cdot C_\delta}\mathbb{F}_n\{f_n(\boldsymbol{T}) > \epsilon\} + \left(\mathbb{E}_{\mathbb{F}_n}\left[e^{2K_g\|\boldsymbol{A}_M\boldsymbol{Z}\|_h}\right]\right)^{1/2}\left(\mathbb{P}_{\mathbb{F}_n}\{\|\boldsymbol{A}_M\boldsymbol{Z}\|_h > C_\delta\}\right)^{1/2} \\ &\leqslant e^{K_g\cdot C_\delta}\mathbb{F}_n\{f_n(\boldsymbol{T}) > \epsilon\} + \left(\mathbb{E}_{\mathbb{F}_n}\left[e^{2K_g\|\boldsymbol{A}_M\boldsymbol{Z}\|_h}\right]\right)^{1/2}\delta^{1/2}.\end{aligned}$$

Using $\lim\sup_{n\to\infty_{\mathbb{F}_n}\in\mathcal{F}_n} \mathbb{F}_n\{f_n(\boldsymbol{T}) > \epsilon\} = 0$ (by the assumption of the lemma) and $\mathbb{E}_{\mathbb{F}_n}\left[e^{2K_g\|\boldsymbol{A}_M\boldsymbol{Z}\|_h}\right] = O(1)$, we get

$$\lim\sup_{n\to\infty_{\mathbb{F}_n}\in\mathcal{F}_n} \mathbb{F}_n^*\{f_n(\boldsymbol{T}) > \epsilon\} = 0.$$

■

4. Bootstrap after selection

This section introduces the bootstrapped test statistic, $\mathcal{P}^*$, for testing $\boldsymbol{A}^T\boldsymbol{\mu}$ after selection. In addition to all the selective CLT assumptions mentioned in Section 2, we also assume that the bootstrapped distribution $\sqrt{n}\boldsymbol{A}^T(\boldsymbol{T}^* - \boldsymbol{T}) = \boldsymbol{A}^T\boldsymbol{Z}^*$ is consistent for $\sqrt{n}\boldsymbol{A}^T(\boldsymbol{T} - \boldsymbol{\mu}) = \boldsymbol{A}^T\boldsymbol{Z}$ before selection, i.e. under $\boldsymbol{S} \sim \mathbb{F}_n$, and we further assume we have a uniformly consistent estimate of the covariance matrix $\boldsymbol{\Sigma}$. Under the assumptions mentioned, we prove that the bootstrapped test statistic $\mathcal{P}^*$ is asymptotically Unif$[0, 1]$.

In order to introduce the bootstrapped test statistic in detail, let us denote as $\widehat{\mathbb{F}}_n$ the bootstrapped distribution of $\boldsymbol{X}$ and let $\widehat{\mathbb{E}}_n$ denote the corresponding expectation with respect to $\widehat{\mathbb{F}}_n$.[14] Denote a bootstrap sample as $\boldsymbol{S}^* = (\boldsymbol{S}_1^*, \ldots, \boldsymbol{S}_n^*)$ with $\boldsymbol{S}_i^* \overset{i.i.d}{\sim} \widehat{\mathbb{F}}_n$, $i = 1, \ldots, n$. For short, denote $\boldsymbol{T}^* = \boldsymbol{T}(\boldsymbol{S}^*)$, the test statistic $\boldsymbol{T}$ computed on the bootstrapped data. Define the bootstrapped selected region as

$$\begin{aligned}&\text{selection}^*(\boldsymbol{T}^*, \boldsymbol{T}, \boldsymbol{\omega}) \\ &:= \left\{\sqrt{n}\boldsymbol{A}_M\widehat{\boldsymbol{T}}_{\boldsymbol{A}} + \boldsymbol{A}_M\widehat{\boldsymbol{C}}\left(\sqrt{n}(\boldsymbol{A}^T\boldsymbol{T}^* - \boldsymbol{A}^T\boldsymbol{T})\right) + \sqrt{n}\boldsymbol{A}_M\widehat{\boldsymbol{C}}\boldsymbol{A}^T\boldsymbol{\mu} + \boldsymbol{\omega} \in \boldsymbol{H}_M\right\} \\ &= \left\{\sqrt{n}\boldsymbol{A}_M\widehat{\boldsymbol{T}}_{\boldsymbol{A}} + \boldsymbol{A}_M\widehat{\boldsymbol{C}}\boldsymbol{A}^T\boldsymbol{Z}^* + \sqrt{n}\boldsymbol{A}_M\widehat{\boldsymbol{C}}\boldsymbol{A}^T\boldsymbol{\mu} + \boldsymbol{\omega} \in \boldsymbol{H}_M\right\},\end{aligned}$$

[14] $\widehat{\mathbb{F}}_n$ can be empirical distribution of $\boldsymbol{S}$ but not necessarily since we can also use wild bootstrap, weighted, etc.

where $\widehat{\boldsymbol{T}}_{\boldsymbol{A}} = (\boldsymbol{I}_p - \widehat{\boldsymbol{C}}\boldsymbol{A}^T)\boldsymbol{T}$ and $\widehat{\boldsymbol{C}} = \widehat{\boldsymbol{\Sigma}}\boldsymbol{A}(\boldsymbol{A}^T\widehat{\boldsymbol{\Sigma}}\boldsymbol{A})^{-1}$ for some estimate $\widehat{\boldsymbol{\Sigma}}$ of $\boldsymbol{\Sigma}$. Given the bootstrapped selection event, the bootstrapped version of the pivot $\mathcal{P}$ is defined as

$$\mathcal{P}^* = \mathcal{P}^*\left(\sqrt{n}\left\|\boldsymbol{A}^T\boldsymbol{T} - \boldsymbol{A}^T\boldsymbol{\mu}\right\|_2, \boldsymbol{T}, \boldsymbol{A}^T\boldsymbol{\mu}\right), \tag{11}$$

where the function $\mathcal{P}^*(t, \boldsymbol{T}, \boldsymbol{A}^T\boldsymbol{\mu})$, $t \in \mathbb{R}$, is defined as a survival function of $\left\|\sqrt{n}\boldsymbol{A}^T(\boldsymbol{T}^* - \boldsymbol{T})\right\|_2$ after bootstrapped selection event and conditional on the data $\boldsymbol{T}$:

$$\begin{aligned}
&\mathcal{P}^*\left(t, \boldsymbol{T}, \boldsymbol{A}^T\boldsymbol{\mu}\right) \\
&:= (\widehat{\mathbb{F}}_n \times \mathbb{G})\left\{\left\|\boldsymbol{A}^T(\sqrt{n}(\boldsymbol{T}^* - \boldsymbol{T}))\right\|_2 \geqslant t \,\middle|\, \text{selection}^*(\boldsymbol{T}^*, \boldsymbol{T}, \boldsymbol{\omega}), \boldsymbol{T}\right\} \\
&= \frac{(\widehat{\mathbb{F}}_n \times \mathbb{G})\left\{\left\|\boldsymbol{A}^T(\sqrt{n}(\boldsymbol{T}^* - \boldsymbol{T}))\right\|_2 \geqslant t, \text{selection}^*(\boldsymbol{T}^*, \boldsymbol{T}, \boldsymbol{\omega}) \,\middle|\, \boldsymbol{T}\right\}}{(\widehat{\mathbb{F}}_n \times \mathbb{G})\left\{\text{selection}^*(\boldsymbol{T}^*, \boldsymbol{T}, \boldsymbol{\omega}) \,\middle|\, \boldsymbol{T}\right\}} \\
&= \frac{\left(\widehat{\mathbb{F}}_n \times \mathbb{G}\right)\left\{\left\|\boldsymbol{A}^T\boldsymbol{Z}^*\right\|_2 \geqslant t, \sqrt{n}\boldsymbol{A}_M\widehat{\boldsymbol{T}}_{\boldsymbol{A}} + \boldsymbol{A}_M\widehat{\boldsymbol{C}}\boldsymbol{A}^T\boldsymbol{Z}^* + \sqrt{n}\boldsymbol{A}_M\widehat{\boldsymbol{C}}\boldsymbol{A}^T\boldsymbol{\mu} + \boldsymbol{\omega} \in \boldsymbol{H}_M \,\middle|\, \boldsymbol{T}\right\}}{\left(\widehat{\mathbb{F}}_n \times \mathbb{G}\right)\left\{\sqrt{n}\boldsymbol{A}_M\widehat{\boldsymbol{T}}_{\boldsymbol{A}} + \boldsymbol{A}_M\widehat{\boldsymbol{C}}\boldsymbol{A}^T\boldsymbol{Z}^* + \sqrt{n}\boldsymbol{A}_M\widehat{\boldsymbol{C}}\boldsymbol{A}^T\boldsymbol{\mu} + \boldsymbol{\omega} \in \boldsymbol{H}_M \,\middle|\, \boldsymbol{T}\right\}}.
\end{aligned}$$

Conditioning on $\boldsymbol{A}^T\boldsymbol{T}^*$, which is equivalent to conditioning on $\boldsymbol{A}^T\boldsymbol{Z}^*$, we get

$$\begin{aligned}
&\mathcal{P}^*\left(t, \boldsymbol{T}, \boldsymbol{A}^T\boldsymbol{\mu}\right) \\
&= \frac{\widehat{\mathbb{E}}_n\left[\mathbb{E}_{\boldsymbol{\omega}\sim\mathbb{G}}\left[\mathbb{I}_{\left\{\|\boldsymbol{A}^T\boldsymbol{Z}^*\|_2 \geqslant t\right\}}\mathbb{I}_{\left\{\boldsymbol{\omega} \in \boldsymbol{H}_M - \sqrt{n}\boldsymbol{A}_M\widehat{\boldsymbol{T}}_{\boldsymbol{A}} - \boldsymbol{A}_M\widehat{\boldsymbol{C}}\boldsymbol{A}^T\boldsymbol{Z}^* - \sqrt{n}\boldsymbol{A}_M\widehat{\boldsymbol{C}}\boldsymbol{A}^T\boldsymbol{\mu}\right\}} \,\middle|\, \boldsymbol{T}, \boldsymbol{A}^T\boldsymbol{T}^*\right]\right]}{\widehat{\mathbb{E}}_n\left[\mathbb{E}_{\boldsymbol{\omega}}\left[\mathbb{I}_{\left\{\boldsymbol{\omega} \in \boldsymbol{H}_M - \sqrt{n}\boldsymbol{A}_M\widehat{\boldsymbol{T}}_{\boldsymbol{A}} - \boldsymbol{A}_M\widehat{\boldsymbol{C}}\boldsymbol{A}^T\boldsymbol{Z}^* - \sqrt{n}\boldsymbol{A}_M\widehat{\boldsymbol{C}}\boldsymbol{A}^T\boldsymbol{\mu}\right\}} \,\middle|\, \boldsymbol{T}, \boldsymbol{A}^T\boldsymbol{T}^*\right]\right]} \\
&= \frac{\widehat{\mathbb{E}}_n\left[\mathbb{I}_{\left\{\|\boldsymbol{A}^T\boldsymbol{Z}^*\|_2 \geqslant t\right\}}\mathbb{G}\left\{\boldsymbol{H}_M - \sqrt{n}\boldsymbol{A}_M\widehat{\boldsymbol{T}}_{\boldsymbol{A}} - \boldsymbol{A}_M\widehat{\boldsymbol{C}}\boldsymbol{A}^T\boldsymbol{Z}^* - \sqrt{n}\boldsymbol{A}_M\widehat{\boldsymbol{C}}\boldsymbol{A}^T\boldsymbol{\mu} \,\middle|\, \boldsymbol{T}, \boldsymbol{A}^T\boldsymbol{T}^*\right\}\right]}{\widehat{\mathbb{E}}_n\left[\mathbb{G}\left\{\boldsymbol{H}_M - \sqrt{n}\boldsymbol{A}_M\widehat{\boldsymbol{T}}_{\boldsymbol{A}} - \boldsymbol{A}_M\widehat{\boldsymbol{C}}\boldsymbol{A}^T\boldsymbol{Z}^* - \sqrt{n}\boldsymbol{A}_M\widehat{\boldsymbol{C}}\boldsymbol{A}^T\boldsymbol{\mu} \,\middle|\, \boldsymbol{T}, \boldsymbol{A}^T\boldsymbol{T}^*\right\}\right]}.
\end{aligned}$$

As mentioned in the simple example (Section 1.4), if we had a (fast and accurate) way of computing probabilities of selection for a given data vector, i.e. $\mathbb{G}\left\{\text{selection}(\boldsymbol{T}, \boldsymbol{\omega})|\boldsymbol{T}\right\}$, as a function of $\boldsymbol{T}$, then we could compute the bootstrap statistic in a usual way using standard sampling with replacement bootstrap. In order to explain how bootstrapped statistic $\mathcal{P}^*$ would be computed, let us denote the bootstrapped samples $\boldsymbol{T}^{*b}$, $b = 1, \ldots, B$. Then the proposed bootstrapped statistic would be computed as the following ratio

$$\frac{\frac{1}{B}\sum_{b=1}^{B} p^*(\boldsymbol{T}, \boldsymbol{T}^{*b})\mathbb{I}_{\left\{\sqrt{n}\|\boldsymbol{A}^T(\boldsymbol{T}^{*b} - \boldsymbol{T})\|_2 \geqslant \sqrt{n}\|\boldsymbol{A}^T(\boldsymbol{T} - \boldsymbol{\mu})\|_2\right\}}}{\frac{1}{B}\sum_{b=1}^{B} p^*(\boldsymbol{T}, \boldsymbol{T}^{*b})},$$

where

$$\begin{aligned}
p^*(\boldsymbol{T}, \boldsymbol{T}^*) &= \mathbb{G}\left\{\text{selection}^*(\boldsymbol{T}^*, \boldsymbol{T}, \boldsymbol{\omega})\middle|\boldsymbol{T}^*, \boldsymbol{T}\right\} \\
&= \mathbb{G}\left\{\boldsymbol{H}_M - \sqrt{n}\boldsymbol{A}_M\widehat{\boldsymbol{T}}_{\boldsymbol{A}} - \boldsymbol{A}_M\widehat{\boldsymbol{C}}\left(\sqrt{n}\boldsymbol{A}^T(\boldsymbol{T}^* - \boldsymbol{T})\right) - \sqrt{n}\boldsymbol{A}_M\widehat{\boldsymbol{C}}\boldsymbol{A}^T\boldsymbol{\mu} \,\middle|\, \boldsymbol{T}, \boldsymbol{T}^*\right\}.
\end{aligned} \tag{12}$$

In the simple example in Section 1.4, we had an explicit form for p^*. However, in more complicated examples it is computationally expensive to get good approximations of p^*, thus we turn to the wild bootstrap and MC techniques to compute this pivot.

The goal is now to show that the $\mathcal{P}^*$ is asymptotically Unif$[0, 1]$ uniformly across $\mathbb{F}_n \in \mathcal{F}_n$. In order to show that, let us introduce some important notation. For $\boldsymbol{\tau} \in \mathbb{R}^a$ and $\boldsymbol{u}_{\boldsymbol{A}} \in \mathbb{R}^p$, define

$$\begin{aligned}
Q_2\left(\boldsymbol{\tau}; \boldsymbol{z}_{\boldsymbol{A}}, \boldsymbol{\Delta}\right) &:= \mathbb{G}\left\{\boldsymbol{\omega} : \boldsymbol{\omega} + \boldsymbol{A}_M\boldsymbol{z}_{\boldsymbol{A}} + \boldsymbol{A}_M\boldsymbol{C}\boldsymbol{\tau} + \boldsymbol{A}_M\boldsymbol{\Delta} \in \boldsymbol{H}_M\right\} \\
&= \mathbb{G}\left\{\boldsymbol{H}_M - \boldsymbol{A}_M\boldsymbol{z}_{\boldsymbol{A}} - \boldsymbol{A}_M\boldsymbol{C}\boldsymbol{\tau} - \boldsymbol{A}_M\boldsymbol{\Delta}\right\} \\
&= \mathbb{G}\left\{\boldsymbol{H}_M - \sqrt{n}\boldsymbol{A}_M\boldsymbol{t}_{\boldsymbol{A}} - \boldsymbol{A}_M(\boldsymbol{C}\boldsymbol{\tau} + \sqrt{n}\boldsymbol{C}\boldsymbol{A}^T\boldsymbol{\mu})\right\},
\end{aligned}$$

where as before $\sqrt{n}\boldsymbol{t}_{\boldsymbol{A}} := \boldsymbol{z}_{\boldsymbol{A}} + \boldsymbol{\Delta} - \sqrt{n}\boldsymbol{C}\boldsymbol{A}^T\boldsymbol{\mu}$, and

$$\begin{aligned} Q_3(\boldsymbol{z}_{\boldsymbol{A}};\boldsymbol{\Delta}) &:= \int_{\mathbb{R}^a} Q_2\left(\boldsymbol{\tau};\boldsymbol{z}_{\boldsymbol{A}},\boldsymbol{\Delta}\right)\frac{1}{(2\pi)^{-a/2}|\boldsymbol{\Sigma}_{\boldsymbol{A}}|^{-1/2}}e^{-\frac{1}{2}\boldsymbol{\tau}^T\boldsymbol{\Sigma}_{\boldsymbol{A}}^{-1}\boldsymbol{\tau}}d\boldsymbol{\tau} \\ &= \int_{\mathbb{R}^a}\mathbb{G}\left\{\boldsymbol{H}_M - \sqrt{n}\boldsymbol{A}_M\boldsymbol{t}_{\boldsymbol{A}} - \boldsymbol{A}_M\left(\boldsymbol{C}\boldsymbol{\tau} + \sqrt{n}\boldsymbol{C}\boldsymbol{A}^T\boldsymbol{\mu}\right)\right\}\frac{e^{-\frac{1}{2}\boldsymbol{\tau}^T\boldsymbol{\Sigma}_{\boldsymbol{A}}^{-1}\boldsymbol{\tau}}}{(2\pi)^{-a/2}|\boldsymbol{\Sigma}_{\boldsymbol{A}}|^{-1/2}}d\boldsymbol{\tau} \\ &= \int_{\mathbb{R}^a}\mathbb{G}\left\{\boldsymbol{H}_M - \sqrt{n}\boldsymbol{A}_M\boldsymbol{t}_{\boldsymbol{A}} - \sqrt{n}\boldsymbol{A}_M\boldsymbol{C}\boldsymbol{s}\right\}\frac{\sqrt{n}}{(2\pi)^{-a/2}|\boldsymbol{\Sigma}_{\boldsymbol{A}}|^{-1/2}}e^{-\frac{n}{2}(\boldsymbol{s}-\boldsymbol{A}^T\boldsymbol{\mu})^T\boldsymbol{\Sigma}_{\boldsymbol{A}}^{-1}(\boldsymbol{s}-\boldsymbol{A}^T\boldsymbol{\mu})}d\boldsymbol{s}, \end{aligned}$$

where $\boldsymbol{\Sigma}_{\boldsymbol{A}} = \boldsymbol{A}^T\boldsymbol{\Sigma}\boldsymbol{A}$ in the last equality we did the change of variables $\boldsymbol{\tau} + \sqrt{n}\boldsymbol{A}^T\boldsymbol{\mu} = \sqrt{n}\boldsymbol{s}$.

In addition to the assumptions on the randomization and the local alternatives assumption, we need the asymptotic result of selective CLT to hold (P), consistency of estimated variance and the consistency of bootstrap pre-selection as stated.

- Uniform consistency of the variance estimate: for any $\delta > 0$, the variance estimate $\widehat{\boldsymbol{\Sigma}}$ of $\boldsymbol{\Sigma}$ satisfies
$$\lim_{n\to\infty}\sup_{\mathbb{F}_n\in\mathcal{F}_n}\mathbb{F}_n\left\{\left\|\widehat{\boldsymbol{\Sigma}} - \boldsymbol{\Sigma}\right\|_2 > \delta\right\} = 0. \tag{Var}$$
This implies
$$\lim_{n\to\infty}\sup_{\mathbb{F}_n\in\mathcal{F}_n}\mathbb{F}_n\left\{\left\|\widehat{\boldsymbol{\Sigma}}_{\boldsymbol{A}} - \boldsymbol{\Sigma}_{\boldsymbol{A}}\right\|_2 > \delta\right\} = 0$$
and
$$\lim_{n\to\infty}\sup_{\mathbb{F}_n\in\mathcal{F}_n}\mathbb{F}_n\left\{\left\|\widehat{\boldsymbol{C}} - \boldsymbol{C}\right\|_2 > \delta\right\} = 0,$$
where $\widehat{\boldsymbol{\Sigma}}_{\boldsymbol{A}}$ and $\widehat{\boldsymbol{C}}$ are the corresponding quantities computed based on $\widehat{\boldsymbol{\Sigma}}$.
- Uniform consistency of bootstrap before selection: for any $\delta > 0$, we have
$$\lim_{n\to\infty}\sup_{\mathbb{F}_n\in\mathcal{F}_n}\mathbb{F}_n\left\{\sup_{\boldsymbol{t}\in\mathbb{R}^a}\left|\widehat{\mathbb{F}}_n\left\{\sqrt{n}\boldsymbol{A}^T(\boldsymbol{T}^* - \boldsymbol{T}) \leqslant \boldsymbol{t}\right\} - \mathbb{P}_{\boldsymbol{G}}\left\{\widehat{\boldsymbol{\Sigma}}_{\boldsymbol{A}}^{1/2}\boldsymbol{G} \leqslant \boldsymbol{t} \,\middle|\, \boldsymbol{T}\right\}\right| > \delta\right\} = 0, \tag{B}$$
where $\boldsymbol{G} \sim \mathcal{N}_a\left(\boldsymbol{0}, \boldsymbol{I}_a\right)$ given $\boldsymbol{T}$, i.e.
$$\sup_{\boldsymbol{t}\in\mathbb{R}^a}\left|\widehat{\mathbb{F}}_n\left\{\sqrt{n}\boldsymbol{A}^T(\boldsymbol{T}^* - \boldsymbol{T}) \leqslant \boldsymbol{t}\right\} - \mathbb{P}_{\boldsymbol{G}}\left\{\widehat{\boldsymbol{\Sigma}}_{\boldsymbol{A}}^{1/2}\boldsymbol{G} \leqslant \boldsymbol{t}|\boldsymbol{T}\right\}\right| = o_{\mathbb{F}_n}(1)$$
uniformly over $\mathbb{F}_n \in \mathcal{F}_n$.

The following theorem proves the bootstrap statistic $\mathcal{P}^*$ is asymptotically distributed as Unif$[0,1]$ under the conditional distribution $\mathbb{F}_n^*$ of the data.

Theorem 10 (Bootstrap after selection) *Under the assumptions* (B), (LA), (Lip), (NA), (P) *and* (Var) *the following holds*

$$\lim_{n\to\infty}\sup_{\mathbb{F}_n\in\mathcal{F}}\sup_{t\in[0,1]}\left|\mathbb{F}_n^*\left\{\mathcal{P}^*\left(\sqrt{n}\boldsymbol{A}^T(\boldsymbol{T} - \boldsymbol{\mu}), \boldsymbol{T}, \boldsymbol{A}^T\boldsymbol{\mu}\right) \leqslant t\right\} - t\right| = 0.$$

Proof Since for any $t \in \mathbb{R}$

$$\begin{aligned} &\mathcal{P}^*\left(t, \boldsymbol{T}, \boldsymbol{A}^T\boldsymbol{\mu}\right) \\ &= \frac{\mathbb{E}_{\boldsymbol{\omega}\sim\mathbb{G}}\left[\widehat{\mathbb{F}}_n\left\{\sqrt{n}\left\|\boldsymbol{A}^T(\boldsymbol{T}^* - \boldsymbol{T})\right\|_2 \geqslant t, \boldsymbol{\omega} \in \boldsymbol{H}_M - \sqrt{n}\boldsymbol{A}_M\widehat{\boldsymbol{T}}_{\boldsymbol{A}} - \boldsymbol{A}_M\widehat{\boldsymbol{C}}\boldsymbol{A}^T\boldsymbol{Z}^* - \sqrt{n}\boldsymbol{A}_M\widehat{\boldsymbol{C}}\boldsymbol{A}^T\boldsymbol{\mu} \,\middle|\, \boldsymbol{T}, \boldsymbol{\omega}\right\}\right]}{\mathbb{E}_{\boldsymbol{\omega}\sim\mathbb{G}}\left[\widehat{\mathbb{F}}_n\left\{\boldsymbol{\omega} \in \boldsymbol{H}_M - \sqrt{n}\boldsymbol{A}_M\widehat{\boldsymbol{T}}_{\boldsymbol{A}} - \boldsymbol{A}_M\widehat{\boldsymbol{C}}\boldsymbol{A}^T\boldsymbol{Z}^* - \sqrt{n}\boldsymbol{A}_M\widehat{\boldsymbol{C}}\boldsymbol{A}^T\boldsymbol{\mu} \,\middle|\, \boldsymbol{T}, \boldsymbol{\omega}\right\}\right]}, \end{aligned}$$

by the bootstrap consistency assumption (B), we can write

$$\mathcal{P}^*\left(t, \boldsymbol{T}, \boldsymbol{A}^T\boldsymbol{\mu}\right) = \frac{\mathbb{E}_{\boldsymbol{\omega}\sim\mathbb{G}}\left[\mathbb{P}_{\boldsymbol{G}}\left\{\|\hat{\boldsymbol{\Sigma}}_{\boldsymbol{A}}^{1/2}\boldsymbol{G}\|_2 \geqslant t, \boldsymbol{\omega}\in \boldsymbol{H}_M - \sqrt{n}\boldsymbol{A}_M\hat{\boldsymbol{T}}_{\boldsymbol{A}} - \boldsymbol{A}_M\hat{\boldsymbol{C}}\hat{\boldsymbol{\Sigma}}_{\boldsymbol{A}}^{1/2}\boldsymbol{G} - \sqrt{n}\boldsymbol{A}_M\hat{\boldsymbol{C}}\boldsymbol{A}^T\boldsymbol{\mu} \;\middle|\; \boldsymbol{T}, \boldsymbol{\omega}\right\}\right] + E_n}{\mathbb{E}_{\boldsymbol{\omega}\sim\mathbb{G}}\left[\mathbb{P}_{\boldsymbol{G}}\left\{\boldsymbol{\omega}\in \boldsymbol{H}_M - \sqrt{n}\boldsymbol{A}_M\hat{\boldsymbol{T}}_{\boldsymbol{A}} - \boldsymbol{A}_M\hat{\boldsymbol{C}}\hat{\boldsymbol{\Sigma}}_{\boldsymbol{A}}^{1/2}\boldsymbol{G} - \sqrt{n}\boldsymbol{A}_M\hat{\boldsymbol{C}}\boldsymbol{A}^T\boldsymbol{\mu} \;\middle|\; \boldsymbol{T}, \boldsymbol{\omega}\right\}\right] + E'_n},$$

where conditional on $\boldsymbol{T}$ we have $\boldsymbol{G} \sim \mathcal{N}_a(\boldsymbol{0}, \boldsymbol{I}_a)$. Furthermore, for the random variables E_n and E'_n we have $E_n, E'_n = o_{\mathbb{F}_n}(1)$ uniformly over $t \in \mathbb{R}$ and over $\mathbb{F}_n \in \mathcal{F}_n$. By the uniform consistency of the variance, assumption (Var), and Lemma 25 (in Appendix D), the following holds for any $t \in \mathbb{R}$

$$\mathcal{P}^*\left(t, \boldsymbol{T}, \boldsymbol{A}^T\boldsymbol{\mu}\right) = \frac{\mathbb{E}_{\boldsymbol{\omega}\sim\mathbb{G}}\left[\mathbb{P}_{\boldsymbol{G}}\left\{\|\boldsymbol{\Sigma}_{\boldsymbol{A}}^{1/2}\boldsymbol{G}\|_2 \geqslant t, \boldsymbol{\omega}\in \boldsymbol{H}_M - \sqrt{n}\boldsymbol{A}_M\boldsymbol{T}_{\boldsymbol{A}} - \boldsymbol{A}_M\boldsymbol{C}\boldsymbol{\Sigma}_{\boldsymbol{A}}^{1/2}\boldsymbol{G} - \sqrt{n}\boldsymbol{A}_M\boldsymbol{C}\boldsymbol{A}^T\boldsymbol{\mu} \;\middle|\; \boldsymbol{T}, \boldsymbol{\omega}\right\}\right] + e_n}{\mathbb{E}_{\boldsymbol{\omega}\sim\mathbb{G}}\left[\mathbb{P}_{\boldsymbol{G}}\left\{\boldsymbol{\omega}\in \boldsymbol{H}_M - \sqrt{n}\boldsymbol{A}_M\boldsymbol{T}_{\boldsymbol{A}} - \boldsymbol{A}_M\boldsymbol{C}\boldsymbol{\Sigma}_{\boldsymbol{A}}^{1/2}\boldsymbol{G} - \sqrt{n}\boldsymbol{A}_M\boldsymbol{C}\boldsymbol{A}^T\boldsymbol{\mu} \;\middle|\; \boldsymbol{T}, \boldsymbol{\omega}\right\}\right] + e'_n},$$

where $e_n, e'_n = o_{\mathbb{F}_n}(1)$ uniformly over $\mathbb{F}_n \in \mathcal{F}$ and over $t \in \mathbb{R}$. By the law of iterated expectation and using the definitions of Q_2 and Q_3, we further have

$$\begin{aligned}
&\mathcal{P}^*(t, \boldsymbol{T}, \boldsymbol{A}^T\boldsymbol{\mu}) \\
&= \frac{\mathbb{E}_{\boldsymbol{G}}\left[\mathbb{I}_{\left\{\|\boldsymbol{\Sigma}_{\boldsymbol{A}}^{1/2}\boldsymbol{G}\|_2 \geqslant t\right\}}\mathbb{G}\left\{\boldsymbol{H}_M - \sqrt{n}\boldsymbol{A}_M\boldsymbol{T}_{\boldsymbol{A}} - \boldsymbol{A}_M\boldsymbol{C}\boldsymbol{\Sigma}_{\boldsymbol{A}}^{1/2}\boldsymbol{G} - \sqrt{n}\boldsymbol{A}_M\boldsymbol{C}\boldsymbol{A}^T\boldsymbol{\mu} \,\middle|\, \boldsymbol{T}, \boldsymbol{G}\right\}\right] + e_n}{\mathbb{E}_{\boldsymbol{G}}\left[\mathbb{G}\left\{\boldsymbol{H}_M - \sqrt{n}\boldsymbol{A}_M\boldsymbol{T}_{\boldsymbol{A}} - \boldsymbol{A}_M\boldsymbol{C}\boldsymbol{\Sigma}_{\boldsymbol{A}}^{1/2}\boldsymbol{G} - \sqrt{n}\boldsymbol{A}_M\boldsymbol{C}\boldsymbol{A}^T\boldsymbol{\mu} \,\middle|\, \boldsymbol{T}, \boldsymbol{G}\right\}\right] + e'_n} \\
&= \frac{\mathbb{E}_{\boldsymbol{G}}\left[\mathbb{I}_{\left\{\|\boldsymbol{\Sigma}_{\boldsymbol{A}}^{1/2}\boldsymbol{G}\|_2 \geqslant t\right\}}Q_2(\boldsymbol{\Sigma}_{\boldsymbol{A}}^{1/2}\boldsymbol{G}; \boldsymbol{Z}_{\boldsymbol{A}}, \boldsymbol{\Delta}) \,\middle|\, \boldsymbol{T}\right] + e_n}{\mathbb{E}_{\boldsymbol{Z}}\left[Q_2(\boldsymbol{\Sigma}_{\boldsymbol{A}}^{1/2}\boldsymbol{G}; \boldsymbol{Z}_{\boldsymbol{A}}, \boldsymbol{\Delta}) \middle| \boldsymbol{T}\right] + e'_n} \\
&= \frac{\mathbb{E}_{\boldsymbol{G}}\left[\mathbb{I}_{\left\{\|\boldsymbol{\Sigma}_{\boldsymbol{A}}^{1/2}\boldsymbol{G}\|_2 \geqslant t\right\}}Q_2(\boldsymbol{\Sigma}_{\boldsymbol{A}}^{1/2}\boldsymbol{G}; \boldsymbol{Z}_{\boldsymbol{A}}, \boldsymbol{\Delta}) \,\middle|\, \boldsymbol{T}\right] + e_n}{Q_3\left(\boldsymbol{Z}_{\boldsymbol{A}}; \boldsymbol{\Delta}\right) + e'_n} \\
&= \frac{\frac{\mathbb{E}_{\boldsymbol{G}}\left[\mathbb{I}_{\left\{\|\boldsymbol{\Sigma}_{\boldsymbol{A}}^{1/2}\boldsymbol{G}\|_2 \geqslant t\right\}}Q_2(\boldsymbol{\Sigma}_{\boldsymbol{A}}^{1/2}\boldsymbol{G}; \boldsymbol{Z}_{\boldsymbol{A}}, \boldsymbol{\Delta}) \,\middle|\, \boldsymbol{T}\right]}{Q_3(\boldsymbol{Z}_{\boldsymbol{A}}; \boldsymbol{\Delta})} + \frac{e_n}{Q_3(\boldsymbol{Z}_{\boldsymbol{A}}; \boldsymbol{\Delta})}}{1 + \frac{e'_n}{Q_3(\boldsymbol{Z}_{\boldsymbol{A}}; \boldsymbol{\Delta})}},
\end{aligned}$$

where in the last equality we divided both numerator and denominator by $Q_3\left(\boldsymbol{Z}_{\boldsymbol{A}}; \boldsymbol{\Delta}\right)$.

For any $\delta > 0$ and $C > 0$ we have

$$\begin{aligned}
\mathbb{F}_n\left\{\left|\frac{e_n}{Q_3\left(\boldsymbol{Z}_{\boldsymbol{A}}; \boldsymbol{\Delta}\right)}\right| \leqslant \delta\right\} &\geqslant \mathbb{F}_n\left\{|e_n| \leqslant \delta C, \;\; Q_3\left(\boldsymbol{Z}_{\boldsymbol{A}}; \boldsymbol{\Delta}\right) \geqslant C\right\} \\
&\geqslant 1 - \mathbb{F}_n\{|e_n| > \delta C\} - \mathbb{F}_n\left\{Q_3\left(\boldsymbol{Z}_{\boldsymbol{A}}; \boldsymbol{\Delta}\right) < C\right\}.
\end{aligned}$$

Since $e_n = o_{\mathbb{F}_n}(1)$ uniformly across $\mathbb{F}_n \in \mathcal{F}$, we have that $\sup_{\mathbb{F}_n} \mathbb{F}_n\{|e_n| > \delta C\} \to 0$ as $n \to \infty$. This is the part where it becomes crucial to have the denominator $Q_3(\boldsymbol{Z}_{\boldsymbol{A}}; \boldsymbol{\Delta})$ lower-bounded in probability. By the Lower bound lemma (Lemma 24 in Appendix D), there exists C to make $\sup_{\mathbb{F}_n \in \mathcal{F}_n} \mathbb{F}_n\left\{Q_3\left(\boldsymbol{Z}_{\boldsymbol{A}}; \boldsymbol{\Delta}\right) < C\right\}$ arbitrarily small for large enough n. This shows

$$\lim_{n\to\infty} \sup_{\mathbb{F}_n \in \mathcal{F}_n} \mathbb{F}_n\left\{\left|\frac{e_n}{Q_3\left(\boldsymbol{Z}_{\boldsymbol{A}}; \boldsymbol{\Delta}\right)}\right| > \delta\right\} = 0.$$

Similarly, we have

$$\lim_{n\to\infty} \sup_{\mathbb{F}_n \in \mathcal{F}_n} \mathbb{F}_n \left\{ \left| \frac{e'_n}{Q_3\left(\boldsymbol{Z}_{\boldsymbol{A}}; \boldsymbol{\Delta}\right)} \right| > \delta \right\} = 0.$$

Thus, we have

$$\mathcal{P}^*\left(t, \boldsymbol{T}, \boldsymbol{A}^T\boldsymbol{\mu}\right) = \frac{\mathbb{E}_{\boldsymbol{G}}\left[\mathbb{I}_{\left\{\|\boldsymbol{\Sigma}_{\boldsymbol{A}}^{1/2}\boldsymbol{G}\|_2 \geqslant t\right\}} Q_2(\boldsymbol{\Sigma}_{\boldsymbol{A}}^{1/2}\boldsymbol{G}; \boldsymbol{Z}_{\boldsymbol{A}}, \boldsymbol{\Delta}) \,\middle|\, \boldsymbol{T}\right]}{Q_3\left(\boldsymbol{Z}_{\boldsymbol{A}}; \boldsymbol{\Delta}\right)} + g_n,$$

where $g_n = o_{\mathbb{F}_n}(1)$ uniformly across $\mathbb{F}_n \in \mathcal{F}$ and over $t \in \mathbb{R}$. By Lemma 9 in Section 3 we also have $g_n = o_{\mathbb{F}_n^*}(1)$, which allows us to make conditional statements. Since

$$\begin{aligned}
&\frac{\mathbb{E}_{\boldsymbol{G}}\left[\mathbb{I}_{\left\{\|\boldsymbol{\Sigma}_{\boldsymbol{A}}^{1/2}\boldsymbol{G}\|_2 \geqslant t\right\}} Q_2(\boldsymbol{\Sigma}_{\boldsymbol{A}}^{1/2}\boldsymbol{G}; \boldsymbol{Z}_{\boldsymbol{A}}, \boldsymbol{\Delta}) \,\middle|\, \boldsymbol{T}\right]}{Q_3\left(\boldsymbol{Z}_{\boldsymbol{A}}; \boldsymbol{\Delta}\right)} \\
&= \frac{\int_{\|\boldsymbol{s}-\boldsymbol{A}^T\boldsymbol{\mu}\|_2 \geqslant t} \mathbb{G}\{\boldsymbol{H}_M - \sqrt{n}\boldsymbol{A}_M\boldsymbol{T}_{\boldsymbol{A}} - \sqrt{n}\boldsymbol{A}_M\boldsymbol{C}\boldsymbol{s}\} \exp\left(-\frac{n}{2}(\boldsymbol{s}-\boldsymbol{A}^T\boldsymbol{\mu})^T\boldsymbol{\Sigma}_{\boldsymbol{A}}^{-1}(\boldsymbol{s}-\boldsymbol{A}^T\boldsymbol{\mu})\right) d\boldsymbol{s}}{\int_{\mathbb{R}^a} \mathbb{G}\{\boldsymbol{H}_M - \sqrt{n}\boldsymbol{A}_M\boldsymbol{T}_{\boldsymbol{A}} - \sqrt{n}\boldsymbol{A}_M\boldsymbol{C}\boldsymbol{s}\} \exp\left(-\frac{n}{2}(\boldsymbol{s}-\boldsymbol{A}^T\boldsymbol{\mu})^T\boldsymbol{\Sigma}_{\boldsymbol{A}}^{-1}(\boldsymbol{s}-\boldsymbol{A}^T\boldsymbol{\mu})\right) d\boldsymbol{s}} \\
&= \mathcal{P}\left(t, \boldsymbol{T}_{\boldsymbol{A}}, \boldsymbol{A}^T\boldsymbol{\mu}\right),
\end{aligned}$$

the survival function of $\|\sqrt{n}\boldsymbol{A}^T(\boldsymbol{T}-\boldsymbol{\mu})\|_2$ when the data is normally distributed, we get that the conclusion holds by the selective CLT assumption.

■

5. Examples for one view of the data

We describe how to compute the plugin CLT and the bootstrapped pivots in specific examples. The bootstrapped test statistic can be computed for the most of the problems in [Tian Harris et al., 2016]; here, we describe two specific examples, both with random design matrix, in detail. The first one selects the model using the randomized Lasso and the second one using marginal screening.

5.1. Lasso with random design

In this section, we discuss the correlation model where the predictor matrix $\boldsymbol{X} \in \mathbb{R}^{n\times p}$, with p columns (predictors), is random and the selection of the relevant subset $E \subset \{1, \ldots, p\}$ of predictors is done using randomized Lasso on $(\boldsymbol{X}, \boldsymbol{y})$, where $\boldsymbol{y} \in \mathbb{R}^n$ is the response vector. Here we provide inference for the population quantities corresponding to model E, i.e. for $\boldsymbol{\beta}_E^* = (\mathbb{E}\left[\boldsymbol{X}_E^T\boldsymbol{X}_E\right])^{-1}\mathbb{E}\left[\boldsymbol{X}_E^T\boldsymbol{y}\right]$, where the expectations are under the pre-selection distribution of $(\boldsymbol{X}, \boldsymbol{y})$. To adjust for selection, the inference is done under the conditional distribution, where the conditioning is on the event that the randomized Lasso applied to $(\boldsymbol{X}, \boldsymbol{y})$ chooses the model E. The selection event that the randomized Lasso selects the predictors in E is asymptotically affine. Finally, testing for the whole vector $\boldsymbol{\beta}_E^*$ or the individual coefficients $\beta_{E,j}^*$, $j = 1, \ldots, |E|$, corresponding to the predictors from E, or any linear combination of $\boldsymbol{\beta}_E^*$, is done using either plugin CLT or the bootstrapped test statistic that is asymptotically uniform under the conditional distribution.

Let us now describe data generating mechanism and the selection event in detail. Assume our data vectors $(\boldsymbol{x}_i, y_i) \in \mathbb{R}^p \times \mathbb{R}$, $i = 1, \ldots, n$, are i.i.d. with a joint distribution $\mathbb{F}_n \in \mathcal{F}_n$. Denote the response vector as $\boldsymbol{y} = (y_1, \ldots, y_n)^T$, the random design matrix as $\boldsymbol{X}$ with rows $\boldsymbol{x}_i^T$, $\boldsymbol{x}_i \in \mathbb{R}^p$:

$$\boldsymbol{y} = \begin{bmatrix} y_1 \\ \vdots \\ y_n \end{bmatrix}_{n\times 1}, \quad \boldsymbol{X} = \begin{bmatrix} \boldsymbol{x}_1^T \\ \vdots \\ \boldsymbol{x}_n^T \end{bmatrix}_{n\times p}.$$

Let us scale the matrix $\boldsymbol{X}$ by $\frac{1}{\sqrt{n}}$ and call the resulting matrix all $\boldsymbol{X}$. The randomized LASSO [Tian and Taylor, 2015b], [Tian Harris et al., 2016] applied to $(\boldsymbol{X}, \boldsymbol{y})$ solves

$$\widehat{\boldsymbol{\beta}}(\boldsymbol{X}, \boldsymbol{y}, \boldsymbol{\omega}) = \underset{\boldsymbol{\beta} \in \mathbb{R}^p}{\operatorname{argmin}} \frac{1}{2} \|\boldsymbol{y} - \boldsymbol{X}\boldsymbol{\beta}\|_2^2 + \lambda \|\boldsymbol{\beta}\|_1 + \frac{\epsilon}{2} \|\boldsymbol{\beta}\|_2^2 - \boldsymbol{\omega}^T \boldsymbol{\beta}, \quad (\boldsymbol{X}, \boldsymbol{y}, \boldsymbol{\omega}) \sim \mathbb{F}_n^n \times \mathbb{G}, \tag{13}$$

where $\boldsymbol{\omega} \sim \mathbb{G}$ represents the noise from a known distribution $\mathbb{G}$ on $\mathbb{R}^p$ and ϵ is a constant of order $O\left(\frac{1}{\sqrt{n}}\right)$.[15] Assume the randomized Lasso selects $(E, \boldsymbol{s}_E)$, where $E \subset \{1, \ldots, p\}$ is the candidate set of variables and $\boldsymbol{s}_E \in \{\pm 1\}^{|E|}$ are their signs. The the solution of (13) becomes $\widehat{\boldsymbol{\beta}} = \widehat{\boldsymbol{\beta}}(\boldsymbol{X}, \boldsymbol{y}, \boldsymbol{\omega}) = \begin{pmatrix} \widehat{\boldsymbol{\beta}}_E \\ \boldsymbol{0} \end{pmatrix}$.

After looking E, we provide inference for the population parameter $\boldsymbol{\beta}_E^*(\mathbb{F}_n)$ where

$$\boldsymbol{\beta}_E^*(\mathbb{F}) = \mathbb{E}_{\mathbb{F}} \left[\boldsymbol{X}_E^T \boldsymbol{X}_E\right]^{-1} \mathbb{E}_{\mathbb{F}} \left[\boldsymbol{X}_E^T \boldsymbol{y}\right].$$

The goal is to provide inference for a linear combination of $\boldsymbol{\beta}_E^*(\mathbb{F}_n)$. In order to get the distribution of a chosen test statistic under the null, we need to sample data form the conditional distribution (conditional on Lasso in (13) selecting model E).

The selection event consists of all such data and randomization pairs such that solving the randomized Lasso (13) for $(\boldsymbol{X}, \boldsymbol{y}, \boldsymbol{\omega})$ gives the same $(E, \boldsymbol{s}_E)$ pair:

$$\widehat{M}_{(E, \boldsymbol{s}_E)} = \left\{(\boldsymbol{X}, \boldsymbol{y}, \boldsymbol{\omega}) : \widehat{\boldsymbol{\beta}}_{-E}(\boldsymbol{X}, \boldsymbol{y}, \boldsymbol{\omega}) = 0, \ \operatorname{sign}(\widehat{\boldsymbol{\beta}}_E(\boldsymbol{X}, \boldsymbol{y}, \boldsymbol{\omega})) = \boldsymbol{s}_E\right\}.$$

Let us denote with $f_{\boldsymbol{x},y}$ and g the densities of distributions $\mathbb{F}_n$ and $\mathbb{G}$, respectively. By the pull-back trick from [Tian Harris et al., 2016], to sample from $(\boldsymbol{X}, \boldsymbol{y}, \boldsymbol{\omega}) \in \widehat{M}_{(E, \boldsymbol{s}_E)}$, it suffices to sample $(\boldsymbol{X}, \boldsymbol{y}, \boldsymbol{\beta}_E, \boldsymbol{u}_{-E})$ from a density proportional to

$$\left(\prod_{i=1}^n f_{\boldsymbol{x},y}(\boldsymbol{x}_i, y_i)\right) \cdot g\left(-\boldsymbol{X}^T \boldsymbol{y} + \begin{pmatrix} \boldsymbol{X}_E^T \boldsymbol{X}_E + \epsilon \boldsymbol{I}_{|E|} \\ \boldsymbol{X}_{-E}^T \boldsymbol{X}_E \end{pmatrix} \boldsymbol{\beta}_E + \lambda \begin{pmatrix} \boldsymbol{s}_E \\ \boldsymbol{u}_{-E} \end{pmatrix}\right) \tag{14}$$

supported on $\mathbb{R}^{n \times p} \times \mathbb{R}^n \times \mathbb{R}_{\boldsymbol{s}_E}^{|E|} \times [-1, 1]^{p-|E|}$, where $\mathbb{R}_{\boldsymbol{s}_E}^{|E|}$ denotes the orthant in $\mathbb{R}^{|E|}$ corresponding to sings $\boldsymbol{s}_E$. Usually we do not know $f_{\boldsymbol{x},y}$ explicitly but for doing the inference on $\boldsymbol{\beta}_E^*$ we sample from a simpler space using the asymptotic distributions.

The only randomness coming from the data in (13) is in the gradient of the loss, which can be expressed in terms of the vector $\boldsymbol{D} = \begin{pmatrix} \bar{\boldsymbol{\beta}}_E \\ \boldsymbol{X}_{-E}^T (\boldsymbol{y} - \boldsymbol{X}_E \bar{\boldsymbol{\beta}}_E) \end{pmatrix}$, where $\bar{\boldsymbol{\beta}}_E = (\boldsymbol{X}_E^T \boldsymbol{X}_E)^{-1} \boldsymbol{X}_E^T \boldsymbol{y}$.[16] Asymptotically, this vector is normally distributed

$$\boldsymbol{D} \to \mathcal{N}_p\left(\boldsymbol{\mu}_{\boldsymbol{D}}, \boldsymbol{\Sigma}_{\boldsymbol{D}}\right), \tag{15}$$

as $n \to \infty$ under $\mathbb{F}_n$. In terms of $\boldsymbol{D}$, the sub-gradient equations from solving (13) become

$$\begin{aligned} \omega(\boldsymbol{D}, \boldsymbol{\beta}_E, \boldsymbol{u}_{-E}) &= -\begin{pmatrix} \boldsymbol{X}_E^T \boldsymbol{X}_E & \boldsymbol{0} \\ \boldsymbol{X}_{-E}^T \boldsymbol{X}_E & \boldsymbol{I}_{p-|E|} \end{pmatrix} \boldsymbol{D} + \begin{pmatrix} \boldsymbol{X}_E^T \boldsymbol{X}_E + \epsilon \boldsymbol{I}_{|E|} \\ \boldsymbol{X}_{-E}^T \boldsymbol{X}_E \end{pmatrix} \boldsymbol{\beta}_E + \lambda \begin{pmatrix} \boldsymbol{s}_E \\ \boldsymbol{u}_{-E} \end{pmatrix} \\ &= \boldsymbol{M}\boldsymbol{D} + \boldsymbol{B}\boldsymbol{\beta}_E + \boldsymbol{U}\boldsymbol{u}_{-E} + \boldsymbol{L}, \end{aligned}$$

[15] The term $\frac{\epsilon}{2}\|\boldsymbol{\beta}\|_1$ is added in the objective to make sure the solution exists. Taking ϵ to be of order $O(\frac{1}{\sqrt{n}})$ corresponds to performing randomized Lasso and taking ϵ to be $O(1)$ corresponds to a randomized elastic net [Zou and Hastie, 2005].

[16] In Appendix E we show $\boldsymbol{D}$ is asymptotically linear and the selection region written in terms of $\boldsymbol{D}$ is asymptotically affine.

with $\text{sign}(\boldsymbol{\beta}_E) = \boldsymbol{s}_E$ and $\|\boldsymbol{u}_{-E}\|_\infty \leqslant 1$, where

$$\boldsymbol{M} = -\begin{pmatrix} \boldsymbol{X}_E^T\boldsymbol{X}_E & \boldsymbol{0} \\ \boldsymbol{X}_{-E}^T\boldsymbol{X}_E & \boldsymbol{I}_{p-|E|} \end{pmatrix}, \ \boldsymbol{B} = \begin{pmatrix} \boldsymbol{X}_E^T\boldsymbol{X}_E + \epsilon\boldsymbol{I}_{|E|} \\ \boldsymbol{X}_{-E}^T\boldsymbol{X}_E \end{pmatrix}, \ \boldsymbol{U} = \begin{pmatrix} \boldsymbol{0} \\ \lambda\boldsymbol{I}_{p-|E|} \end{pmatrix}, \ \boldsymbol{L} = \begin{pmatrix} \lambda\boldsymbol{s}_E \\ \boldsymbol{0} \end{pmatrix}.$$

Hence, instead of sampling from the density in (14), we alternatively sample $(\boldsymbol{D}, \boldsymbol{\beta}_E, \boldsymbol{u}_E)$ using the asymptotic normality of $\boldsymbol{D}$.

When testing the hypothesis $H_0 : \boldsymbol{A}^T\boldsymbol{\beta}_E^* = \boldsymbol{\theta}$, we use $\|\boldsymbol{T} - \boldsymbol{\theta}\|_2^2$, where $\boldsymbol{T} = \boldsymbol{A}^T\bar{\boldsymbol{\beta}}_E$ denotes the target, as the test statistic. In general, we can do inference for any parameter $\theta(\mathbb{F}_n) = \boldsymbol{\theta}$, assuming the target $\boldsymbol{T}$ for $\boldsymbol{\theta}$ and the data vector $\boldsymbol{D}$ satisfy the pre-selection CLT:

$$\begin{pmatrix} \boldsymbol{T} \\ \boldsymbol{D} \end{pmatrix} \to \mathcal{N}_{a+p}\left(\begin{pmatrix} \boldsymbol{\theta} \\ \boldsymbol{\mu}_{\boldsymbol{D}} \end{pmatrix}, \begin{pmatrix} \boldsymbol{\Sigma}_{\boldsymbol{T}} & \boldsymbol{\Sigma}_{\boldsymbol{T},\boldsymbol{D}} \\ \boldsymbol{\Sigma}_{\boldsymbol{D},\boldsymbol{T}} & \boldsymbol{\Sigma}_{\boldsymbol{D}} \end{pmatrix}\right)$$

as $n \to \infty$. To get the distribution of the target after selection under the null it suffices to sample $(\boldsymbol{T}, \boldsymbol{\beta}_E, \boldsymbol{u}_{-E})$ since we treat part of the data

$$\boldsymbol{F} = \boldsymbol{D} - \hat{\boldsymbol{\Sigma}}_{\boldsymbol{D},\boldsymbol{T}}\hat{\boldsymbol{\Sigma}}_{\boldsymbol{T}}^{-1}\boldsymbol{T},$$

as fixed, where $\hat{\boldsymbol{\Sigma}}_{\boldsymbol{D},\boldsymbol{T}}$ and $\hat{\boldsymbol{\Sigma}}_{\boldsymbol{T}}$ are the estimates of the respective covariances.

The plugin CLT sampling density of $(\boldsymbol{T}, \boldsymbol{\beta}_E, \boldsymbol{u}_{-E})$ is proportional to

$$\phi_a\left(\hat{\boldsymbol{\Sigma}}_{\boldsymbol{T}}^{-1/2}(\boldsymbol{T} - \boldsymbol{\theta})\right) \cdot g\left(\widetilde{\boldsymbol{M}}\boldsymbol{T} + \boldsymbol{B}\boldsymbol{\beta}_E + \boldsymbol{U}\boldsymbol{u}_{-E} + \widetilde{\boldsymbol{L}}\right), \tag{16}$$

supported on $\mathbb{R}^p \times \mathbb{R}_{\boldsymbol{s}_E}^{|E|} \times [-1,1]^{p-|E|}$, where $\mathbb{R}_{\boldsymbol{s}_E}^{|E|}$ is the orthant in $\mathbb{R}^{|E|}$ corresponding to the signs $\boldsymbol{s}_E$ and

$$\widetilde{\boldsymbol{M}} = \boldsymbol{M}\hat{\boldsymbol{\Sigma}}_{\boldsymbol{D},\boldsymbol{T}}\hat{\boldsymbol{\Sigma}}_{\boldsymbol{T}}^{-1}, \ \widetilde{\boldsymbol{L}} = \boldsymbol{L} + \boldsymbol{M}\boldsymbol{F}.$$

Here ϕ_a is the density of standard multivariate normal in a dimensions and $\hat{\boldsymbol{\Sigma}}_{\boldsymbol{T}}$ is an estimate of $\boldsymbol{\Sigma}_{\boldsymbol{T}}$. We estimate $\boldsymbol{\Sigma}_{\boldsymbol{D},\boldsymbol{T}}$ and $\boldsymbol{\Sigma}_{\boldsymbol{T}}$ non-parametrically using pairs bootstrap [Freedman et al., 1981], [Buja et al., 2014].

The wild bootstrap approximates the distribution of $\bar{\boldsymbol{\beta}}_E - \boldsymbol{\beta}_E^* = (\boldsymbol{X}_E^T\boldsymbol{X}_E)^{-1}\boldsymbol{X}_E^T\boldsymbol{\varepsilon}$ with

$$(\boldsymbol{X}_E^T\boldsymbol{X}_E)^{-1}\boldsymbol{X}_E^T\text{diag}(\hat{\boldsymbol{\varepsilon}})\boldsymbol{\alpha},$$

where $\hat{\boldsymbol{\varepsilon}} = \boldsymbol{y} - \boldsymbol{X}_E\bar{\boldsymbol{\beta}}_E$ are the observed residuals from fitting OLS with response $\boldsymbol{y}$ and the predictors $\boldsymbol{X}_E$ and $\boldsymbol{\alpha} \in \mathbb{R}^n$ are bootstrap weights. We take $\alpha_i \overset{i.i.d.}{\sim} \mathbb{F}_\alpha$ with density f_α and support $\text{supp}(\mathbb{F}_\alpha)$. Hence in the sampling density we replace $\boldsymbol{T}$ with $\boldsymbol{T}(\boldsymbol{\alpha}) + \boldsymbol{\theta}$, where

$$\boldsymbol{T}(\boldsymbol{\alpha}) = \boldsymbol{A}^T(\boldsymbol{X}_E^T\boldsymbol{X}_E)^{-1}\boldsymbol{X}_E^T\text{diag}(\hat{\boldsymbol{\varepsilon}})\boldsymbol{\alpha}.$$

The randomization reconstruction map in the bootstrap case becomes

$$\omega^b(\boldsymbol{\alpha}, \boldsymbol{\beta}_E, \boldsymbol{u}_{-E}) = \widetilde{\boldsymbol{M}}\boldsymbol{T}(\boldsymbol{\alpha}) + \boldsymbol{B}\boldsymbol{\beta}_E + \boldsymbol{U}\boldsymbol{u}_{-E} + \widetilde{\boldsymbol{L}} + \widetilde{\boldsymbol{M}}\boldsymbol{\theta}. \tag{17}$$

In this case the bootstrap density on $(\boldsymbol{\alpha}, \boldsymbol{\beta}_E, \boldsymbol{u}_{-E})$ is proportional to

$$\left(\prod_{i=1}^n f_\alpha(\alpha_i)\right) \cdot g\left(\omega^b(\boldsymbol{\alpha}, \boldsymbol{\beta}_E, \boldsymbol{u}_{-E})\right) \tag{18}$$

and supported on $(\text{supp}(\mathbb{F}_\alpha))^n \times \mathbb{R}_{\boldsymbol{s}_E}^{|E|} \times [-1,1]^{p-|E|}$.

Remark 11 *Note that here we do not assume the selected linear model $\boldsymbol{y} \sim \boldsymbol{X}_E$ is true, i.e. we do not assume $\mathbb{E}[\boldsymbol{\varepsilon}|\boldsymbol{X}] = 0$, where $\boldsymbol{\varepsilon} = \boldsymbol{y} - \boldsymbol{X}_E\boldsymbol{\beta}_E^*$ are the true residuals. If we further assume this holds, we could condition on less. In this case, $\boldsymbol{D}$ would have asymptotic mean $\begin{pmatrix} \boldsymbol{\beta}_E^* \\ \boldsymbol{0} \end{pmatrix}$, so we could condition on $\boldsymbol{D}_E - \boldsymbol{\Sigma}_{\boldsymbol{D}_E,\boldsymbol{T}}\boldsymbol{\Sigma}_{\boldsymbol{T}}^{-1}\boldsymbol{T}$ and marginalize over the null statistic $\boldsymbol{D}_{-E}$.*

5.2. Marginal screening

Take as before $\boldsymbol{X} \in \mathbb{R}^{n\times p}$ to be a random design matrix with columns $\boldsymbol{X}_j \in \mathbb{R}^n$, scaled by $\frac{1}{\sqrt{n}}$, $j = 1, \ldots, p$, and the response vector $\boldsymbol{y} \in \mathbb{R}^n$. Marginal screening computes marginal t-statistics

$$\boldsymbol{S}_j = \frac{\boldsymbol{X}_j^T \boldsymbol{y}}{\hat{\sigma}_j},\ j = 1, \ldots, p,$$

and thresholds their absolute value at some threshold, perhaps $z_{1-\alpha/2}$ where α is some nominal p-value threshold.

Selective inference in the nonrandomized setting for this problem was considered in [Lee and Taylor, 2014] and the randomized one in [Tian Harris et al., 2016]. We compute the bootstrapped test statistic in the randomized setting (the same as [Tian Harris et al., 2016])

$$\widehat{\boldsymbol{\eta}}(\boldsymbol{S}, \boldsymbol{\omega}) = \underset{\boldsymbol{\eta}\in\mathbb{R}^p, \|\boldsymbol{\eta}\|_\infty \leqslant c}{\operatorname{argmin}} \frac{1}{2}\|\boldsymbol{\eta} - \boldsymbol{S}\|_2^2 - \boldsymbol{\omega}^T\boldsymbol{\eta},$$

with $\boldsymbol{\omega} \sim \mathbb{G}$ independent of the data. Conditioning on the set E achieving the threshold c and their signs to be $\boldsymbol{s}_E$, we see that this event is

$$\widehat{M}_{(E,\boldsymbol{s}_E)} = \{(\boldsymbol{S}, \boldsymbol{\omega}) : \widehat{\boldsymbol{\eta}}_E(\boldsymbol{S}, \boldsymbol{\omega}) = c \cdot \boldsymbol{s}_E, \|\widehat{\boldsymbol{\eta}}_{-E}(\boldsymbol{S}, \boldsymbol{\omega})\| < c\}$$

or in terms of the data vector T and optimization variables the selection event becomes

$$\{(\boldsymbol{S}, \boldsymbol{\eta}, \boldsymbol{z}) : \boldsymbol{\eta}_E = c \cdot \boldsymbol{s}_E, \operatorname{diag}(\boldsymbol{s}_E)\boldsymbol{z}_E \geqslant 0, \|\boldsymbol{\eta}_{-E}\|_\infty < c, \boldsymbol{z}_{-E} = 0\}.^{17} \tag{19}$$

Let us describe the joint density of the data and optimization variables, $(\boldsymbol{\eta}_{-E}, \boldsymbol{z}_E)$, assuming the goal is inference for a linear combination of the parameter $\boldsymbol{\beta}_E^*$. Denoting the OLS estimator of using only the predictors in E as $\bar{\boldsymbol{\beta}}_E$, we write the vector $\boldsymbol{S}$ in terms of

$$\boldsymbol{D} = \begin{pmatrix} \boldsymbol{D}_E \\ \boldsymbol{D}_{-E} \end{pmatrix} = \begin{pmatrix} \bar{\boldsymbol{\beta}}_E \\ \boldsymbol{X}_{-E}^T(\boldsymbol{y} - \boldsymbol{X}_E\bar{\boldsymbol{\beta}}_E) \end{pmatrix}, \tag{20}$$

as

$$\boldsymbol{S} = \operatorname{diag}(\mathbf{1}/\hat{\boldsymbol{\sigma}})\boldsymbol{X}^T\boldsymbol{y} = \operatorname{diag}(\mathbf{1}/\hat{\boldsymbol{\sigma}}) \begin{pmatrix} \boldsymbol{X}_E^T\boldsymbol{X}_E & \mathbf{0}_{|E|\times(p-|E|)} \\ \boldsymbol{X}_{-E}^T\boldsymbol{X}_E & \boldsymbol{I}_{p-|E|} \end{pmatrix} \boldsymbol{D} =: -\boldsymbol{M}\boldsymbol{D}, \tag{21}$$

where $\operatorname{diag}(\mathbf{1}/\hat{\boldsymbol{\sigma}})$ is a $p \times p$ matrix with diagonal elements $\hat{\sigma}_j$, $j = 1, \ldots, p$, and zeros elsewhere. The randomization reconstruction map becomes

$$\omega^g(\boldsymbol{D}, \boldsymbol{\eta}_{-E}, \boldsymbol{z}_E) = \begin{pmatrix} c \cdot \boldsymbol{s}_E \\ \boldsymbol{\eta}_{-E} \end{pmatrix} + \boldsymbol{M}\boldsymbol{D} + \begin{pmatrix} \boldsymbol{z}_E \\ \mathbf{0} \end{pmatrix}.$$

Using the asymptotic distribution

$$\boldsymbol{D} \to \mathcal{N}_p(\boldsymbol{\mu}_{\boldsymbol{D}}, \boldsymbol{\Sigma}_{\boldsymbol{D}}),$$

as $n \to \infty$, we sample $(\boldsymbol{D}, \boldsymbol{\eta}_{-E}, \boldsymbol{z}_E)$ from a density proportional to

$$\phi_p\left(\boldsymbol{\Sigma}_{\boldsymbol{D}}^{-1/2}(\boldsymbol{D} - \boldsymbol{\mu}_{\boldsymbol{D}})\right) \cdot g\left(\begin{pmatrix} c \cdot \boldsymbol{s}_E \\ \boldsymbol{\eta}_{-E} \end{pmatrix} + \boldsymbol{M}\boldsymbol{D} + \begin{pmatrix} \boldsymbol{z}_E \\ \mathbf{0} \end{pmatrix}\right),$$

and supported on $\mathbb{R}^p \times [-c, c]^{p-|E|} \times \mathbb{R}^{|E|}_{\boldsymbol{s}_E}$, where we plug in the estimate of $\boldsymbol{\Sigma}_{\boldsymbol{D}}$.

[17] $\boldsymbol{z}$ is the subgradient of the characteristic function corresponding to the set $\{\boldsymbol{\eta} \in \mathbb{R}^p : \|\boldsymbol{\eta}\|_\infty \leqslant c\}$.

If we are testing $H_0 : \boldsymbol{A}^T\boldsymbol{\beta}^*_E = \boldsymbol{\theta}$, for some matrix $\boldsymbol{A} \in \mathbb{R}^{a\times p}$, we sample only a part of the data, i.e. the target $\boldsymbol{T} = \boldsymbol{A}^T\bar{\boldsymbol{\beta}}_E$ here by using the decomposition of $\boldsymbol{D}$ in terms of T and the fixed part

$$\boldsymbol{F} = \boldsymbol{D} - \widehat{\boldsymbol{\Sigma}}_{\boldsymbol{D},\boldsymbol{T}}\widehat{\boldsymbol{\Sigma}}^{-1}_{\boldsymbol{T}}\boldsymbol{T},$$

where $\widehat{\boldsymbol{\Sigma}}_{\boldsymbol{D},\boldsymbol{T}}$ is the estimated covariance between the data vector $\boldsymbol{D}$ and the target and $\widehat{\boldsymbol{\Sigma}}_{\boldsymbol{T}}$ is the estimated variance of $\boldsymbol{T}$. Then we sample $(\boldsymbol{T}, \boldsymbol{\eta}_{-E}, \boldsymbol{z}_E)$ from

$$\phi_a\left(\widehat{\boldsymbol{\Sigma}}^{-1/2}_{\boldsymbol{T}}(\boldsymbol{T}-\boldsymbol{\theta})\right)\cdot g\left(\begin{pmatrix} c\cdot \boldsymbol{s}_E \\ \boldsymbol{\eta}_{-E}\end{pmatrix} + \boldsymbol{M}\boldsymbol{F} + \boldsymbol{M}\widehat{\boldsymbol{\Sigma}}_{\boldsymbol{D},\boldsymbol{T}}\widehat{\boldsymbol{\Sigma}}^{-1}_{\boldsymbol{T}}\boldsymbol{T} + \begin{pmatrix}\boldsymbol{z}_E \\ \boldsymbol{0}\end{pmatrix}\right),$$

with the support $\mathbb{R}^a \times [-c, c]^{p-|E|} \times \mathbb{R}^{|E|}_{\boldsymbol{s}_E}$.

Using the wild bootstrap, we approximate the distribution of $\boldsymbol{X}^T_E\boldsymbol{\varepsilon}$ with $\boldsymbol{X}^T_E\text{diag}(\hat{\varepsilon})\boldsymbol{\alpha}$, i.e. replace $\boldsymbol{T} = \boldsymbol{A}^T(\boldsymbol{X}^T_E\boldsymbol{X}_E)^{-1}\boldsymbol{X}^T_E\boldsymbol{\varepsilon}$ with $\boldsymbol{T}(\boldsymbol{\alpha}) + \boldsymbol{\theta}$, where

$$\boldsymbol{T}(\boldsymbol{\alpha}) = \boldsymbol{A}^T(\boldsymbol{X}^T_E\boldsymbol{X}_E)^T\boldsymbol{X}^T_E\text{diag}(\hat{\boldsymbol{\epsilon}})\boldsymbol{\alpha}.$$

Then the bootstrapped randomization reconstruction map denote the randomization reconstruction map as

$$\omega^b(\boldsymbol{\alpha}, \boldsymbol{\eta}_{-E}, \boldsymbol{z}_E) = \begin{pmatrix} c\cdot \boldsymbol{s}_E \\ \boldsymbol{\eta}_{-E}\end{pmatrix} + \boldsymbol{M}\boldsymbol{F} + \boldsymbol{M}\widehat{\boldsymbol{\Sigma}}_{\boldsymbol{D},\boldsymbol{T}}\widehat{\boldsymbol{\Sigma}}^{-1}_{\boldsymbol{T}}\boldsymbol{T}(\boldsymbol{\alpha}) + \boldsymbol{M}\widehat{\boldsymbol{\Sigma}}_{\boldsymbol{D},\boldsymbol{T}}\widehat{\boldsymbol{\Sigma}}^{-1}_{\boldsymbol{T}}\boldsymbol{\theta} + \begin{pmatrix}\boldsymbol{z}_E \\ \boldsymbol{0}\end{pmatrix}.$$

Now we sample $(\boldsymbol{\alpha}, \boldsymbol{\eta}_{-E}, \boldsymbol{z}_E)$ from the density proportional to

$$\left(\prod_{i=1}^n f_\alpha(\alpha_i)\right)\cdot g(\omega^b(\boldsymbol{\alpha}, \boldsymbol{\eta}_{-E}, \boldsymbol{z}_E))$$

and supported on $\text{supp}(\mathbb{F}_\alpha)^n \times [-c, c]^{p-|E|} \times \mathbb{R}^{|E|}_{\boldsymbol{s}_E}$.

6. Multiple views of the data

In this section, we present a way to do inference after running several randomized model selection procedures and choosing a parameter of interest upon looking at the outcomes of all of them. We present a general sampling framework, which we then illustrate with two examples: randomized forward-stepwise and multiple runs of randomized ℓ_1-penalized logistic regression.

Given data $\boldsymbol{S} \sim \mathbb{F}$ at step n, assume we solve K randomized model selection procedures (views) of the following form

$$\underset{\boldsymbol{\beta}}{\text{minimize}}\ \mathcal{F}_k(\boldsymbol{\beta}, \boldsymbol{S}) - \boldsymbol{\omega}^T_k\boldsymbol{\beta}, \quad (\boldsymbol{S}, \boldsymbol{\omega}_k) \sim \mathbb{F}\times\mathbb{G}_k, \tag{22}$$

where $\mathcal{F}_k$ is some function, e.g.

$$\mathcal{F}_k(\boldsymbol{\beta}, \boldsymbol{S}) = \ell_k(\boldsymbol{\beta}; \boldsymbol{S}) + \mathcal{P}_k(\boldsymbol{\beta}) + \frac{\epsilon_k}{2}\|\boldsymbol{\beta}\|^2_2,$$

for some loss function ℓ_k, penalty $\mathcal{P}_k$ and $\epsilon_k > 0$ is added to make sure the solution exists. Here, $\boldsymbol{\omega}_k$ is the randomization distributed from $\mathbb{G}_k$ with density g_k. Each view defines a set of KKT conditions from solving (22), i.e. from solving the randomization reconstruction equation

$$\boldsymbol{\omega}_k = \partial_{\boldsymbol{\beta}}\mathcal{F}_k(\boldsymbol{\beta}, \boldsymbol{S}),$$

which we write in terms of the optimization variables, denoted as $\boldsymbol{V}_k \in \mathbb{R}^{n_k}$ and a variable $\boldsymbol{D}_k = D_k(\boldsymbol{S}) \in \mathbb{R}^{p_k}$ that is only a function of the data $\boldsymbol{S}$. Hence, we write the KKT conditions for k-th view are as

$$\boldsymbol{\omega}_k = \omega_k(\boldsymbol{D}_k, \boldsymbol{V}_k)$$

with the constraint $\boldsymbol{V}_k \in \mathcal{V}_k \subset \mathbb{R}^{n_k}$.

Now in order to do inference post-selection, we need the marginal distribution of the data $\boldsymbol{S}$ conditional on the event that running all the K procedures for $\boldsymbol{S} \sim \mathbb{F}$ and $(\boldsymbol{\omega}_1, \ldots, \boldsymbol{\omega}_K) \sim \prod_{k=1}^K \mathbb{G}_k$ gives the same sequence of models $(E_1, \ldots, E_K)$. To get that distribution, we sample $\boldsymbol{S}$ and the optimization variables $(\boldsymbol{V}_1, \ldots, \boldsymbol{V}_K)$ from the density

$$f(\boldsymbol{S}) \cdot \prod_{k=1}^{K} g_k\left(\omega_k(D_k(\boldsymbol{S}), \boldsymbol{V}_k)\right) \tag{23}$$

restricted to $V_k \in \mathcal{V}_k$, $k = 1, \ldots K$. Depending on the inferential goal, we do not necessarily sample the whole data $\boldsymbol{S}$ but a function of it.

Given the models selected at each of the K views, an analyst decides to test a parameter $H_0 : \theta(\mathbb{F}) = \boldsymbol{\theta} \in \mathbb{R}^a$ using an estimator $\boldsymbol{T} = T(\boldsymbol{S})$. We now describe how to simplify the density in (23) so that we sample $\boldsymbol{T}$ instead of $\boldsymbol{S}$. We call $\boldsymbol{T}$ the target and assume that pre-selection there is a CLT

$$\begin{pmatrix} \boldsymbol{T} \\ \boldsymbol{D}_k \end{pmatrix} \to \mathcal{N}_{a+p_k}\left(\begin{pmatrix} \boldsymbol{\theta} \\ \boldsymbol{\mu}_k \end{pmatrix}, \begin{pmatrix} \boldsymbol{\Sigma}_{\boldsymbol{T}} & \boldsymbol{\Sigma}_{\boldsymbol{T},\boldsymbol{D}_k} \\ \boldsymbol{\Sigma}_{\boldsymbol{D}_k,\boldsymbol{T}} & \boldsymbol{\Sigma}_{\boldsymbol{D}_k} \end{pmatrix}\right)$$

as $n \to \infty$ for all $k = 1, \ldots, K$. Denoting the estimates of $\boldsymbol{\Sigma}_{\boldsymbol{D}_k,\boldsymbol{T}} = \mathrm{Cov}(\boldsymbol{D}_k, \boldsymbol{T})$ and $\boldsymbol{\Sigma}_{\boldsymbol{T}} = \mathrm{Var}(\boldsymbol{T})$ as $\widehat{\boldsymbol{\Sigma}}_{\boldsymbol{D}_k,\boldsymbol{T}}$ and $\widehat{\boldsymbol{\Sigma}}_{\boldsymbol{T}}$, we treat

$$\boldsymbol{F}_k = \boldsymbol{D}_k - \widehat{\boldsymbol{\Sigma}}_{\boldsymbol{D}_k,\boldsymbol{T}} \widehat{\boldsymbol{\Sigma}}_{\boldsymbol{T}}^{-1} \boldsymbol{T}$$

as fixed in the sampler. Hence, we write the randomization reconstruction in at step k only as a function of $\boldsymbol{T}$ and $(\boldsymbol{V}_1, \ldots, \boldsymbol{V}_K)$ and the plugin CLT sampling density of these variables becomes

$$\phi_a\left(\widehat{\boldsymbol{\Sigma}}_{\boldsymbol{T}}^{-1/2}(\boldsymbol{T} - \boldsymbol{\theta})\right) \cdot \prod_{i=1}^{K} \left(g_k(\omega_k(\boldsymbol{F}_k + \widehat{\boldsymbol{\Sigma}}_{\boldsymbol{D}_k,\boldsymbol{T}} \widehat{\boldsymbol{\Sigma}}_{\boldsymbol{T}}^{-1} \boldsymbol{T}, \boldsymbol{V}_k)) \cdot J_k(\boldsymbol{V}_k)\right)$$

with the constrains $(\boldsymbol{V}_1, \ldots, \boldsymbol{V}_K) \in \prod_{k=1}^K \mathcal{V}_k$, where J_k is the Jacobian coming from the change of density at k-th view. Here, ϕ_a is the density of a standard Gaussian in a dimensions.

Using wild bootstrap, $\boldsymbol{T}$ gets replaced by $T(\boldsymbol{\alpha}) + \boldsymbol{\theta}$, where $T(\boldsymbol{\alpha})$ is consistent for $\mathcal{N}(\boldsymbol{0}, \boldsymbol{\Sigma}_{\boldsymbol{T}})$. Hence the sampling density of the bootstrapped weights $\boldsymbol{\alpha} \in \mathbb{R}^n$ and the optimization variables $\boldsymbol{V}_1, \ldots, \boldsymbol{V}_K$ becomes

$$\left(\prod_{i=1}^{n} f_\alpha(\alpha_i)\right) \times \prod_{k=1}^{K} \left(g_k\left(\omega_k(\boldsymbol{F}_k + \widehat{\boldsymbol{\Sigma}}_{\boldsymbol{D}_k,\boldsymbol{T}} \widehat{\boldsymbol{\Sigma}}_{\boldsymbol{T}}^{-1}(\boldsymbol{T}(\boldsymbol{\alpha}) + \boldsymbol{\theta}), \boldsymbol{V}_k)\right) \cdot J_k(\boldsymbol{V}_k)\right).$$

with the constraints on the bootstrap weights $\boldsymbol{\alpha} \in \mathrm{supp}(\mathbb{F}_\alpha)^n$ and the optimization variables as before $(\boldsymbol{V}_1, \ldots, \boldsymbol{V}_K) \in \prod_{k=1}^K \mathcal{V}_k$.

6.1. Forward stepwise

We revise the randomized forward stepwise example from [Tian Harris et al., 2016], which is a special case of a Kac-Rice test [Taylor et al., 2013], and describe how the sampling here works once we specify the target of inference.

Our data consists of the response vector $\boldsymbol{y} \in \mathbb{R}^n$ and the design matrix $\boldsymbol{X} \in \mathbb{R}^{n \times p}$ with rows $\boldsymbol{x}_i^T$, $i = 1, \ldots, n$, where $(\boldsymbol{x}_i, y_i) \overset{i.i.d.}{\sim} \mathbb{F}_n$, $i = 1, \ldots, n$, with density $f_{\boldsymbol{x},y}$. Let us scale $\boldsymbol{X}$ by $\frac{1}{\sqrt{n}}$. We perform forward stepwise algorithm with the total number of steps fixed at K to choose K predictors.

In the K steps of forward stepwise, the selection event is characterized by a sequence of indices $\boldsymbol{j} = (j_1, \ldots, j_K)$ with their corresponding signs $\boldsymbol{s} = (s_1, \ldots, s_K)$ that enter the model in that particular

order, forming the active set at step K. Denote the active set at step k as $\mathcal{A}_k = \{j_1, j_2, \ldots, j_k\}$ for all $k = 1, \ldots, K$. At the k-th step the randomized forward stepwise solves the following program

$$\underset{\boldsymbol{\eta} \in \mathcal{B}_k}{\text{maximize}}\, \boldsymbol{\eta}^T \left(\boldsymbol{X}_{-\mathcal{A}_{k-1}}^T \mathcal{P}_{\mathcal{A}_{k-1}}^{\perp} \boldsymbol{y} + \boldsymbol{\omega}_k \right), \tag{24}$$

where $(\boldsymbol{X}, \boldsymbol{y}) \times \boldsymbol{\omega}_k \sim \mathbb{F}_n^n \times \mathbb{G}_k$. Here, $\{\boldsymbol{\omega}_k\}_{k=1}^K$ is a sequence of independent randomization variables with $\boldsymbol{\omega}_k$ coming from a given distribution $\mathbb{G}_k$ in $\mathbb{R}^{p-k+1}$ with support $\text{supp}(\mathbb{G}_k)$ and density g_k. $\boldsymbol{X}_{-\mathcal{A}_k}$ are the columns of $\boldsymbol{X}$ except for the ones corresponding to the current active set $\mathcal{A}_k$, and

$$\mathcal{B}_k = \{\boldsymbol{\eta} \in \mathbb{R}^{p-k+1} : \|\boldsymbol{\eta}\|_1 \leqslant 1\}.$$

Let us denote the solution of (24) at the k-th step as $\hat{\boldsymbol{\eta}}_k$. The selection event of interest here is given by conditioning on the sign and index on nonzero coordinate of $\hat{\boldsymbol{\eta}}_k$ for each $k = 1, \ldots, K$. Hence, we want to sample from the density of the data and the randomization conditional on this selection event. The randomization reconstruction map for the k-th step, from the subgradient equation is given by

$$\omega_k(\boldsymbol{y}, \boldsymbol{z}_k) = -\boldsymbol{X}_{-\mathcal{A}_{k-1}}^T \mathcal{P}_{\mathcal{A}_{k-1}}^{\perp} \boldsymbol{y} + \boldsymbol{z}_k,$$

where, subdifferential $\boldsymbol{z}_k \in \mathbb{R}^{p-k+1}$ from the k-th step is restricted to the normal cone $\boldsymbol{z}_k \in \partial I_{\mathcal{B}_k}(\hat{\boldsymbol{\eta}}_k)$.[18]

As in [Tian Harris et al., 2016], in order to sample $((\boldsymbol{X}, \boldsymbol{y}), \boldsymbol{\omega}_1, \ldots, \boldsymbol{\omega}_K)$ from the selective density, we sample $((\boldsymbol{X}, \boldsymbol{y}), \boldsymbol{z}_1, \ldots, \boldsymbol{z}_K)$ from the density proportional to

$$\left(\prod_{i=1}^n f_{\boldsymbol{x},y}(\boldsymbol{x}_i, y_i) \right) \cdot \prod_{k=1}^K g_k \left(\boldsymbol{z}_k - \boldsymbol{X}_{-\mathcal{A}_{k-1}}^T \mathcal{P}_{\mathcal{A}_{k-1}}^{\perp} \boldsymbol{y} \right), \tag{25}$$

supported on $\mathbb{R}^{n \times p} \times \mathbb{R}^n \times \prod_{k=1}^K \partial I_{\mathcal{B}_k}(\hat{\boldsymbol{\eta}}_k)$.

6.1.1. Inference after FS

After doing K steps of forward stepwise, an analyst looks at the sequence $\{\mathcal{A}_k\}_{k=1}^K$ and chooses model E in whichever way she wants. Now the goal is to inference for the population OLS parameters $\boldsymbol{\beta}_E^*$ in the model $\boldsymbol{y} \sim \boldsymbol{X}_E$ and here we present a sampling scheme for testing a given parameter.

As in the Lasso example, we simplify the sampling above since we are interested in testing a particular linear combination of the parameters $\boldsymbol{\beta}_E^*$ or any parameter $\boldsymbol{\theta} = \theta(\mathbb{F}_n)$ in general. First note that $\boldsymbol{X}_{-\mathcal{A}_{k-1}}^T \mathcal{P}_{\mathcal{A}_{k-1}}^{\perp} \boldsymbol{y}$ can be expressed as $\boldsymbol{Q}_k \boldsymbol{X}^T \boldsymbol{y}$, where

$$\boldsymbol{Q}_k = \left[\boldsymbol{X}_{-\mathcal{A}_{k-1}}^T \boldsymbol{X}_{\mathcal{A}_{k-1}} (\boldsymbol{X}_{\mathcal{A}_{k-1}}^T \boldsymbol{X}_{\mathcal{A}_{k-1}})^{-1} \quad -\boldsymbol{I}_{p-(k-1)} \right]. \tag{26}$$

Now we write the sampling density in terms of a new data vector

$$\boldsymbol{D} = \begin{pmatrix} \bar{\boldsymbol{\beta}}_E \\ \boldsymbol{N}_{-E} \end{pmatrix} = \begin{pmatrix} \bar{\boldsymbol{\beta}}_E \\ \boldsymbol{X}_{-E}^T (\boldsymbol{y} - \boldsymbol{X}_E \bar{\boldsymbol{\beta}}_E) \end{pmatrix}, \tag{27}$$

where $\bar{\boldsymbol{\beta}}_E$ is the least squares estimate of the response $\boldsymbol{y}$ and covariate matrix $\boldsymbol{X}_E$. Since as n tends to infinity,

$$\boldsymbol{D} \to \mathcal{N}_p \left(\boldsymbol{\mu}_{\boldsymbol{D}}, \boldsymbol{\Sigma}_{\boldsymbol{D}} \right), \tag{28}$$

we use the asymptotic normality of $\boldsymbol{D}$ to sample $(\boldsymbol{D}, \boldsymbol{z}_1, \ldots, \boldsymbol{z}_K)$ from the density proportional to

$$\phi_p \left(\boldsymbol{\Sigma}_{\boldsymbol{D}}^{-1/2} (\boldsymbol{D} - \boldsymbol{\mu}_{\boldsymbol{D}}) \right) \cdot \prod_{k=1}^K g_k \left(\boldsymbol{z}_k + \boldsymbol{M}_k \boldsymbol{D} \right),$$

[18] The details are given in [Tian Harris et al., 2016].

and supported on $\mathbb{R}^p \times \prod_{k=1}^K \partial I_{\mathcal{B}_k}(\hat{\boldsymbol{\eta}}_k)$. Here, $\boldsymbol{M}_k = -\boldsymbol{Q}_k \begin{pmatrix} \boldsymbol{X}_E^T \boldsymbol{X}_E & \boldsymbol{0} \\ \boldsymbol{X}_{-E}^T \boldsymbol{X}_E & \boldsymbol{I}_{p-|E|} \end{pmatrix}$ and we use an estimate of $\boldsymbol{\Sigma_D}$ in the sampler.

If we want to test $H_0 : \boldsymbol{A}^T \boldsymbol{\beta}_E^* = \boldsymbol{\theta}$, we set the target to be

$$\boldsymbol{T} = \boldsymbol{A}^T \bar{\boldsymbol{\beta}}_{\mathcal{A}_K},$$

but in general we do inference for any $\boldsymbol{\theta} = \theta(\mathbb{F})$ as long as there is a joint CLT for $\boldsymbol{T}$ and $\boldsymbol{D}$ preselection. Denote $\boldsymbol{\Sigma}_{\boldsymbol{D},\boldsymbol{T}}$ to be the covariance between the data vector $\boldsymbol{D}$ and the target $\boldsymbol{T}$ and $\boldsymbol{\Sigma}_{\boldsymbol{T}}$ the variance of $\boldsymbol{T}$ and as $\hat{\boldsymbol{\Sigma}}_{\boldsymbol{D},\boldsymbol{T}}$ and $\hat{\boldsymbol{\Sigma}}_{\boldsymbol{T}}$ their respective estimates. Denote the part of the data we keep fixed as

$$\boldsymbol{F} = \boldsymbol{D} - \hat{\boldsymbol{\Sigma}}_{\boldsymbol{D},\boldsymbol{T}} \hat{\boldsymbol{\Sigma}}_{\boldsymbol{T}}^{-1} \boldsymbol{T}.$$

Hence we sample $(\boldsymbol{T}, \boldsymbol{z}_1, \ldots, \boldsymbol{z}_K)$ from the density proportional to

$$\phi_a \left(\boldsymbol{\Sigma}_{\boldsymbol{T}}^{-1/2} (\boldsymbol{T} - \boldsymbol{\theta}) \right) \cdot \prod_{k=1}^K g_k \left(\boldsymbol{z}_k + \boldsymbol{M}_k \boldsymbol{F} + \widetilde{\boldsymbol{M}}_k \boldsymbol{T} \right) \tag{29}$$

and supported on $\mathbb{R}^K \times \prod_{k=1}^K \partial I_{\mathcal{B}_k}(\hat{\boldsymbol{\eta}}_k)$, where $\widetilde{\boldsymbol{M}}_k = \boldsymbol{M}_k \hat{\boldsymbol{\Sigma}}_{\boldsymbol{D},\boldsymbol{T}} \hat{\boldsymbol{\Sigma}}_{\boldsymbol{T}}^{-1}$.

Wild bootstrap approximates the distribution of $\bar{\boldsymbol{\beta}}_E - \boldsymbol{\beta}_E^* = (\boldsymbol{X}_E^T \boldsymbol{X}_E)^{-1} \boldsymbol{X}_E^T \boldsymbol{\varepsilon}$, where $\boldsymbol{\varepsilon} = \boldsymbol{y} - \boldsymbol{X}_E \boldsymbol{\beta}_E^*$ are the true residuals, with

$$(\boldsymbol{X}_E^T \boldsymbol{X}_E)^{-1} \boldsymbol{X}_E^T \mathrm{diag}(\hat{\boldsymbol{\varepsilon}}) \boldsymbol{\alpha},$$

where $\hat{\boldsymbol{\varepsilon}} = \boldsymbol{y} - \boldsymbol{X}_E \bar{\boldsymbol{\beta}}_E$ are the observed residuals after fitting the OLS with the response $\boldsymbol{y}$ and covariates $\boldsymbol{X}_E$. Hence we replace the target $\boldsymbol{T}$ with $\boldsymbol{T}(\boldsymbol{\alpha}) + \boldsymbol{\theta}$, where $\boldsymbol{T}(\boldsymbol{\alpha})$ is the bootstrapped target

$$\boldsymbol{T}(\boldsymbol{\alpha}) = \boldsymbol{A}^T (\boldsymbol{X}_E^T \boldsymbol{X}_E)^{-1} \boldsymbol{X}_E^T \mathrm{diag}(\hat{\boldsymbol{\varepsilon}}) \boldsymbol{\alpha}.$$

In this case, the randomization reconstruction map at k-th step becomes

$$\omega_k^b(\boldsymbol{\alpha}, \boldsymbol{z}_k) = \boldsymbol{z}_k + \boldsymbol{M}_k \boldsymbol{F} + \widetilde{\boldsymbol{M}}_k \boldsymbol{\theta} + \widetilde{\boldsymbol{M}}_k \boldsymbol{T}(\boldsymbol{\alpha}) = \boldsymbol{z}_k + \boldsymbol{M}_k \boldsymbol{F} + \widetilde{\boldsymbol{M}}_k \boldsymbol{\theta} + \overline{\boldsymbol{M}}_k \boldsymbol{\alpha}, \tag{30}$$

where $\overline{\boldsymbol{M}}_k = \widetilde{\boldsymbol{M}}_k \boldsymbol{A}^T (\boldsymbol{X}_E^T \boldsymbol{X}_E)^{-1} \boldsymbol{X}_E^T \mathrm{diag}(\hat{\boldsymbol{\varepsilon}})$, hence linear in the bootstrap weights $\boldsymbol{\alpha}$ and the subgradient $\boldsymbol{z}_k$. Taking into account all the randomization reconstruction maps for $k = 1, \ldots, K$, the bootstrap density of $(\boldsymbol{\alpha}, \boldsymbol{z}_1, \ldots, \boldsymbol{z}_K)$ is proportional to

$$\left(\prod_{i=1}^n f_\alpha(\alpha_i) \right) \cdot \prod_{k=1}^K g_k \left(\omega_k^b(\boldsymbol{\alpha}, \boldsymbol{z}_k) \right) \tag{31}$$

and supported on $(\mathrm{supp}(\mathbb{F}_\alpha))^n \times \prod_{i=1}^n \partial I_{\mathcal{B}_k}(\hat{\boldsymbol{\eta}}_k)$.

6.2. Multiple views for GLM loss and group Lasso penalty

We are given the response vector $\boldsymbol{y} \in \mathbb{R}^n$ and the covariance matrix $\boldsymbol{X} \in \mathbb{R}^{n \times p}$ with the rows denoted as $\boldsymbol{x}_i \in \mathbb{R}^p$, $i = 1, \ldots, n$. Assume $(\boldsymbol{x}_i, \boldsymbol{y}_i) \overset{i.i.d.}{\sim} \mathbb{F}$, $i = 1, \ldots, n$. Given a sequence of randomization variables $\{\boldsymbol{\omega}_k\}_{k=1}^K$, each distributed from $\mathbb{G}_k$ on $\mathbb{R}^p$ and independent of $(\boldsymbol{X}, \boldsymbol{y})$, we do a total of K model selection procedures (or views or queries of the data), each solving the following optimization problem

$$\hat{\boldsymbol{\beta}}_k = \underset{\boldsymbol{\beta} \in \mathbb{R}^p}{\arg\min}\, \ell_k(\boldsymbol{\beta}; (\boldsymbol{X}, \boldsymbol{y})) - \boldsymbol{\omega}_k^T \boldsymbol{\beta} + \mathcal{P}_k(\boldsymbol{\beta}) + \frac{\epsilon_k}{2} \|\boldsymbol{\beta}\|_2^2, \quad (\boldsymbol{X}, \boldsymbol{y}) \times \boldsymbol{\omega}_k \sim \mathbb{F}^n \times \mathbb{G}_k, \tag{32}$$

where ℓ_k is a loss function and $\mathcal{P}_k$ a penalty term for k-th view, $k = 1, \ldots, K$. For the logistic loss function the above loss becomes

$$\ell_k(\boldsymbol{\beta}; (\boldsymbol{X}, \boldsymbol{y})) = -\sum_{i=1}^{n} \left(y_i \log \pi(\boldsymbol{x}_i^T \boldsymbol{\beta}) + (1 - y_i) \log(1 - \pi(\boldsymbol{x}_i^T \boldsymbol{\beta})) \right),$$

where $\pi(x) = \frac{e^x}{1+e^x}$, and the group Lasso penalty above becomes

$$\mathcal{P}_k(\boldsymbol{\beta}) = \sum_{g \in G_k} \lambda_{k,g} \|\boldsymbol{\beta}_g\|_2,$$

where G_k defines the partition $\cup_{g \in G_k} g = \{1, \ldots, p\}$ and $\lambda_{k,g}$ are the group weights.[19] This example has been done in detail in [Tian Harris et al., 2016] for Gaussian loss and one view of the data. Here, we revise their selective plugin CLT density before presenting the bootstrap in this case.

Let us introduce some more notation before writing KKT conditions. Denote with A_k the set of active groups selected by solving the above procedure and with E_k the set of active variables selected, i.e. $E_k = \cup_{g \in A_k} g$. Hence, we write $\hat{\boldsymbol{\beta}}_k = \begin{pmatrix} \hat{\boldsymbol{\beta}}_{k,E_k} \\ \mathbf{0} \end{pmatrix}$, where in the interchangeable notation $(\hat{\boldsymbol{\beta}}_{k,g})_{g \in A_k} = \hat{\boldsymbol{\beta}}_{k,E_k}$. Now let $\bar{\boldsymbol{\beta}}_{E_k} \in \mathbb{R}^{|E_k|}$ denote the solution of the unpenalized and non-randomized version of (32) including only the variables in E_k ($\boldsymbol{X}_{E_k}$ with rows $\boldsymbol{x}_{i,E_k}$, $i = 1, \ldots, n$), i.e.

$$\bar{\boldsymbol{\beta}}_{E_k} = \underset{\boldsymbol{\beta} \in \mathbb{R}^{|E_k|}}{\arg\min} -\sum_{i=1}^{n} \left(y_i \log \pi(\boldsymbol{x}_{i,E_k}^T \boldsymbol{\beta}) + (1 - y_i) \log(1 - \pi(\boldsymbol{x}_{i,E_k}^T \boldsymbol{\beta})) \right),$$

hence $\bar{\boldsymbol{\beta}}_{E_k}$ satisfies $\boldsymbol{X}_{E_k}^T (\boldsymbol{y} - \pi(\boldsymbol{X}_{E_k} \bar{\boldsymbol{\beta}}_{E_k})) = 0$.[20] Define the data vector

$$\boldsymbol{D}_k = \begin{pmatrix} \bar{\boldsymbol{\beta}}_{E_k} \\ \boldsymbol{N}_{-E_k} \end{pmatrix}.$$

Let us also write the following quantities

$$\begin{aligned}
\pi_{E_k}(\boldsymbol{\beta}) &= \frac{\exp(\boldsymbol{X}_{E_k} \boldsymbol{\beta})}{1 + \exp(\boldsymbol{X}_{E_k} \boldsymbol{\beta})} \\
W_{E_k}(\boldsymbol{\beta}) &= \operatorname{diag}\left(\pi_{E_k}(\boldsymbol{\beta})(1 - \pi_{E_k}(\boldsymbol{\beta}))\right) \\
Q_{E_k}(\boldsymbol{\beta}) &= \boldsymbol{X}_{E_k}^T W_{E_k}(\boldsymbol{\beta}) \boldsymbol{X}_{E_k} \\
C_{E_k}(\boldsymbol{\beta}) &= \boldsymbol{X}_{-E_k}^T W_{E_k}(\boldsymbol{\beta}) \boldsymbol{X}_{E_k} \\
I_{E_k}(\boldsymbol{\beta}) &= C_{E_k}(\boldsymbol{\beta}) Q_{E_k}^{-1}(\boldsymbol{\beta})
\end{aligned}$$

for any $\boldsymbol{\beta} \in \mathbb{R}^{|E_k|}$.

Using the Taylor expansion of the gradient of the loss (the details for one query are in [Tian Harris et al., 2016]), the subgradient equation (32) becomes

$$\boldsymbol{\omega}_k = -\begin{pmatrix} Q_{E_k}(\bar{\boldsymbol{\beta}}_{E_k}) & \mathbf{0} \\ C_{E_k}(\bar{\boldsymbol{\beta}}_{E_k}) & \boldsymbol{I}_{p-|E_k|} \end{pmatrix} \boldsymbol{D}_k + \begin{pmatrix} Q_{E_k}(\bar{\boldsymbol{\beta}}_{E_k}) + \epsilon \boldsymbol{I}_{|E_k|} \\ C_{E_k}(\bar{\boldsymbol{\beta}}_{E_k}) \end{pmatrix} \hat{\boldsymbol{\beta}}_{k,E_k} + \begin{pmatrix} \left(\lambda_{k,g} \frac{\hat{\boldsymbol{\beta}}_{k,g}}{\|\hat{\boldsymbol{\beta}}_{k,g}\|_2} \right)_{g \in A_k} \\ (\boldsymbol{z}_{k,h})_{h \in -A_k} \end{pmatrix}$$

with the constrains $\|\boldsymbol{z}_{k,h}\|_2 \leqslant \lambda_{k,h}, \quad \forall\, h \in -A_k$. Conditioning on the active directions, denoted as

$$\boldsymbol{u}_{k,g} = \frac{\hat{\boldsymbol{\beta}}_{k,g}}{\|\hat{\boldsymbol{\beta}}_{k,g}\|_2}, \quad g \in A_k,$$

[19] We use g to denote both a group and its elements.

[20] A function π is applied component-wise to a vector.

we write $\hat{\boldsymbol{\beta}}_{k,g} = \gamma_{k,g}\boldsymbol{u}_{k,g}$, where $\gamma_{k,g} = \|\hat{\boldsymbol{\beta}}_{k,g}\|_2$. Now, we write the subgradient equation at view k as an affine function in terms of the data vector $\boldsymbol{D}_k$ and the optimization variables $(\gamma_{k,g})_{g\in A_k}$, $\gamma_{k,g} \in \mathbb{R}$, and $(\boldsymbol{z}_{k,h})_{h\in -E_k}$, $\boldsymbol{z}_{k,h} \in \mathbb{R}^{|h|}$, restricted to

$$\gamma_{k,g} > 0, \quad \forall g \in A_k, \quad \text{and} \quad \|\boldsymbol{z}_{k,h}\|_2 \leqslant \lambda_{k,h}, \quad \forall h \in -A_k. \tag{33}$$

In a compact form, we write the randomization reconstruction as

$$\boldsymbol{\omega}_k = \boldsymbol{M}_k\boldsymbol{D}_k + \boldsymbol{\Gamma}_k(\gamma_{k,g})_{g\in A_k} + \boldsymbol{Z}_k(\boldsymbol{z}_{k,h})_{h\in -E_k} + \boldsymbol{L}_k,$$

where $\boldsymbol{M}_k \in \mathbb{R}^{p\times p}$, $B_k \in \mathbb{R}^{p\times|E_k|}$ and $\boldsymbol{Z}_k \in \mathbb{R}^{p\times(p-|E_k|)}$ denote the matrices

$$\begin{aligned}
\boldsymbol{M}_k &= -\begin{pmatrix} Q_{E_k}(\bar{\boldsymbol{\beta}}_{E_k}) & \boldsymbol{0} \\ C_{E_k}(\bar{\boldsymbol{\beta}}_{E_k}) & \boldsymbol{I}_{p-|E_k|} \end{pmatrix}, \\
\boldsymbol{\Gamma}_k &= \begin{pmatrix} Q_{E_k}(\bar{\boldsymbol{\beta}}_{E_k}) + \epsilon\boldsymbol{I}_{|E_k|} \\ C_{E_k}(\bar{\boldsymbol{\beta}}_{E_k}) \end{pmatrix} \mathrm{diag}((\boldsymbol{u}_{k,g})_{g\in A_k}), \\
\boldsymbol{Z}_k &= \begin{pmatrix} \boldsymbol{0} \\ \boldsymbol{I}_{p-|E_k|} \end{pmatrix}
\end{aligned}$$

and the vector $\boldsymbol{L}_k \in \mathbb{R}^p$ denotes

$$\boldsymbol{L}_k = \begin{pmatrix} (\lambda_{k,g}\boldsymbol{u}_{k,g})_{g\in A_k} \\ \boldsymbol{0} \end{pmatrix}.$$

After looking at the outcomes of all K views, i.e. the sets $E_1, \ldots, E_K$, an analyst decides on a model E. It could be that the analyst decides on $E = \cup_{k=1}^K E_k$ the union of all active variables across the views but not necessarily. Assume the analysis is interested in testing

$$\boldsymbol{\beta}_E^* = \left(\mathbb{E}_{\mathbb{F}}\left[\boldsymbol{X}_E^T\boldsymbol{X}_E\right]\right)^{-1}\mathbb{E}_{\mathbb{F}}\left[\boldsymbol{X}_E^T\boldsymbol{y}\right]$$

or some linear combination of it. Also, let $\bar{\boldsymbol{\beta}}_E$ be non-penalized and non-randomized minimizer of the loss using only the covariates in $\boldsymbol{X}_E$.

If we are interested in testing $H_0 : \boldsymbol{A}^T\boldsymbol{\beta}_E^* = \boldsymbol{\theta}$ for some matrix $\boldsymbol{A} \in \mathbb{R}^{|E|\times a}$ chosen after looking at the active sets selected across all the views, then we use $\|\boldsymbol{A}^T(\bar{\boldsymbol{\beta}}_E - \boldsymbol{\beta}_E^*)\|_2$ as the test statistic and use the samples of $\boldsymbol{A}^T\bar{\boldsymbol{\beta}}_E$ after selection to get the distribution of the test statistics under the null. Denote the target variable as

$$\boldsymbol{T} = \boldsymbol{A}^T\bar{\boldsymbol{\beta}}_E.$$

Focusing on the data part pre-selection, there is a CLT for all $k = 1, \ldots, K$:

$$\begin{pmatrix} \boldsymbol{T} \\ \boldsymbol{D}_k \end{pmatrix} \xrightarrow{d} \mathcal{N}_{a+p}\left(\begin{pmatrix} \boldsymbol{\theta} \\ \boldsymbol{\mu}_k \end{pmatrix}, \boldsymbol{\Sigma}_k\right) \tag{34}$$

as $n \to \infty$, for some $\boldsymbol{\mu}_k \in \mathbb{R}^p$ and $\boldsymbol{\Sigma}_k \in \mathbb{R}^{(a+p)\times(a+p)}$. Since the only data part we want to sample is $\boldsymbol{T}$, we condition on the rest, i.e. the part of the data independent of $\boldsymbol{T}$ if (34) would be exact, i.e.

$$\boldsymbol{D}_k - \boldsymbol{\Sigma}_{\boldsymbol{D}_k,\boldsymbol{T}}\boldsymbol{\Sigma}_{\boldsymbol{T}}^{-1}\boldsymbol{T}, \tag{35}$$

where

$$\boldsymbol{\Sigma}_{\boldsymbol{D}_k,\boldsymbol{T}} = \mathrm{Cov}\left(\boldsymbol{D}_k, \boldsymbol{T}\right) \quad \text{and} \quad \boldsymbol{\Sigma}_{\boldsymbol{T}} = \mathrm{Var}(\boldsymbol{T}).$$

For the two covariances matrices, we use their parametric or non-parametric (pairs bootstrap) estimates, denotes as $\widehat{\boldsymbol{\Sigma}}_{\boldsymbol{D}_k,\boldsymbol{T}}$ and $\widehat{\boldsymbol{\Sigma}}_{\boldsymbol{T}}$ respectively, and denote the quantity corresponding to (35) computed based on the covariance estimates as $\boldsymbol{F}_k$.

The randomization reconstruction at view k becomes

$$\boldsymbol{\omega}_k = \widetilde{\boldsymbol{M}}_k\boldsymbol{T} + \boldsymbol{\Gamma}_k(\gamma_{k,g})_{g\in A_k} + \boldsymbol{Z}_k(\boldsymbol{z}_{k,h})_{h\in -A_k} + \widetilde{\boldsymbol{L}}_k, \tag{36}$$

where

$$\widetilde{\boldsymbol{M}}_k = \boldsymbol{M}_k \widehat{\boldsymbol{\Sigma}}_{\boldsymbol{D}_k,\boldsymbol{T}} \widehat{\boldsymbol{\Sigma}}_{\boldsymbol{T}}^{-1}, \quad \widetilde{\boldsymbol{L}}_k = \boldsymbol{L}_k + \boldsymbol{M}_k \boldsymbol{F}_k$$

and the optimization variables $(\gamma_{k,g})_{g\in A_k}$ and $(\boldsymbol{z}_{k,h})_{h\in -A_k}$ are restricted as in (33). The plugin CLT selective density in terms of the target $\boldsymbol{T}$ and the optimization variables $(\gamma_{k,g})_{g\in A_k}$ and $(\boldsymbol{z}_{k,h})_{h\in -A_k}$, $k = 1, \ldots, K$, is proportional to

$$\phi_a\left(\widehat{\boldsymbol{\Sigma}}_{\boldsymbol{T}}^{-1/2}(\boldsymbol{T}-\boldsymbol{\theta})\right) \cdot \prod_{k=1}^{K}\left(g_k(\widetilde{\boldsymbol{M}}_k\boldsymbol{T} + \boldsymbol{\Gamma}_k(\gamma_{k,g})_{g\in A_k} + \boldsymbol{Z}_k(\boldsymbol{z}_{k,h})_{h\in -A_k} + \widetilde{\boldsymbol{L}}_k) \cdot J_k\right) \tag{37}$$

and restricted to (33), where ϕ_a is the density of a standard Gaussian, $\mathcal{N}_a(\boldsymbol{0}, \boldsymbol{I}_a)$, in a dimensions, and J_k is the Jacobian coming from the change of density at view k.

6.2.1. Wild bootstrap

Using the Taylor expansion of

$$\boldsymbol{X}_E^T(\boldsymbol{y} - \pi_E(\bar{\boldsymbol{\beta}}_E)) = \boldsymbol{0}$$

around $\boldsymbol{\beta}_E^*$, we have

$$0 \approx \boldsymbol{X}_E^T(\boldsymbol{y} - \pi_E(\boldsymbol{\beta}_E^*)) - \boldsymbol{X}_E^T\boldsymbol{W}_E(\boldsymbol{\beta}_E^*)\boldsymbol{X}_E(\bar{\boldsymbol{\beta}}_E - \boldsymbol{\beta}_E^*) = 0.$$

This gives

$$\bar{\boldsymbol{\beta}}_E - \boldsymbol{\beta}_E^* = Q_E(\boldsymbol{\beta}_E^*)^{-1}\sum_{i=1}^{n}\boldsymbol{x}_{E,i}^T(y_i - \pi(\boldsymbol{x}_{E,i}^T\boldsymbol{\beta}_E^*)) = Q_E(\boldsymbol{\beta}_E^*)^{-1}\boldsymbol{X}_E^T\boldsymbol{\varepsilon}, \tag{38}$$

where $\boldsymbol{x}_{E,i}^T$, $i = 1, \ldots, n$, are the rows of $\boldsymbol{X}_E$ and $\boldsymbol{\varepsilon} = \boldsymbol{y} - \pi_E(\boldsymbol{\beta}_E^*)$ are the true residuals. Let $\hat{\boldsymbol{\varepsilon}}$ be the observed residuals after fitting unpenalized model with covariates E, i.e. $\hat{\boldsymbol{\varepsilon}} = \boldsymbol{y} - \pi_E(\bar{\boldsymbol{\beta}}_E)$. Using wild bootstrap,

$$Q_E(\boldsymbol{\beta}_E^*)^{-1}\boldsymbol{X}_E^T\mathrm{diag}(\hat{\boldsymbol{\varepsilon}})\boldsymbol{\alpha},$$

where $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_n) \in \mathbb{R}^n$ with $\alpha_i \overset{i.i.d.}{\sim} \mathbb{F}_\alpha$, approximates the distribution of (38). Denote

$$T(\boldsymbol{\alpha}) = \boldsymbol{A}^T Q_E(\boldsymbol{\beta}_E^*)^{-1}\boldsymbol{X}_E^T\mathrm{diag}(\hat{\boldsymbol{\varepsilon}}).$$

Using wild bootstrap, $\boldsymbol{T}$ gets replaced with $T(\boldsymbol{\alpha}) + \boldsymbol{\theta}$. The sampling density on the bootstrap weights $\boldsymbol{\alpha} \in \mathbb{R}^n$ and the optimization variables $(\gamma_{k,g})_{g\in A_k}$ and $(\boldsymbol{z}_{k,g})_{g\in -A_k}$, $k = 1, \ldots, K$, becomes

$$\left(\prod_{i=1}^{n} f_\alpha(\alpha_i)\right) \cdot \left(\prod_{k=1}^{K} g_k\left(\omega_k(\overline{\boldsymbol{M}}_k\boldsymbol{\alpha} + \boldsymbol{F}_k(\gamma_{k,g})_{g\in A_k} + \boldsymbol{Z}_k(\boldsymbol{z}_{k,g})_{h\in -A_k}) + \bar{\boldsymbol{L}}_k\right) \cdot J_k\right), \tag{39}$$

where

$$\overline{\boldsymbol{M}}_k = \widetilde{\boldsymbol{M}}_k\boldsymbol{A}^T Q_E(\boldsymbol{\beta}_E^*)^{-1}\boldsymbol{X}_E^T\mathrm{diag}(\hat{\boldsymbol{\varepsilon}}), \quad \bar{\boldsymbol{L}}_k = \widetilde{\boldsymbol{L}}_k + \widetilde{\boldsymbol{M}}\boldsymbol{\theta}, \tag{40}$$

with the weights restricted to $\boldsymbol{\alpha} \in \mathrm{supp}(\mathbb{F}_\alpha)^n$ and the optimization variables as in (33),

7. Data splitting revisited

Classical data splitting uses a random subsample of the data to choose the model and the other part of the data to do inference assuming the model selected at the first stage is correct. Hence, a part of the data (called split data here) is used only for model selection. Here, we also select the model based on a part of the data but we use the whole data for the inference by conditioning only on the model selected in the first stage. Hence, while the classical data splitting conditions on all

of the first stage data, we condition only on the model selected using the first stage data, doing the inference using the leftover data together with the leftover information in the split data (after using it to select the model). This idea, called data carving, of reusing the leftover information from the first stage has been introduced in [Fithian et al., 2014]. Here, we provide the computational tools to do data carving efficiently.

Let us now describe the setup and the procedure in detail. Given data matrix $\boldsymbol{S} \in \mathbb{R}^{n\times(p+1)}$, we resample m of its rows randomly without replacement and denote the random split data as $\boldsymbol{S}_1 \in \mathbb{R}^{m\times(p+1)}$, where $m \leqslant n$ is given. Assume pre-selection $\boldsymbol{S} \sim \mathbb{F}^n$, i.e. the rows are sampled i.i.d. from $\mathbb{F}$. Denote the full loss as $\ell(\boldsymbol{\beta}; \boldsymbol{S})$ and the split loss as $\ell_1(\boldsymbol{\beta}; \boldsymbol{S}_1)$. In the standard data splitting the model is selected by solving an optimization problem on the split data and then inference is done only on the leftover part of the data that has not been used in the model selection stage. An optimization problem of the form

$$\underset{\boldsymbol{\beta}}{\text{minimize}}\ \frac{1}{\rho}\ell_1(\boldsymbol{\beta}; \boldsymbol{S}_1) + \mathcal{P}(\boldsymbol{\beta}), \tag{41}$$

where $\rho = \frac{m}{n}$ and $\mathcal{P}(\boldsymbol{\beta})$ is a penalty, is solved at the model selection stage and a model E is selected. Then the inference is done using $\boldsymbol{S}_2 \in \mathbb{R}^{(n-m)\times(p+1)}$ assuming E is correct, where $\boldsymbol{S}_2$ consists of the rows of $\boldsymbol{S}$ not used to form $\boldsymbol{S}_1$. Denoting

$$\delta(\boldsymbol{\beta}; \boldsymbol{S}) = \rho\ell(\boldsymbol{\beta}; \boldsymbol{S}) - \ell_1(\boldsymbol{\beta}; \boldsymbol{S}_1),$$

(41) is equivalent to solving

$$\underset{\boldsymbol{\beta}}{\text{minimize}}\ \ell(\boldsymbol{\beta}; \boldsymbol{S}) - \frac{1}{\rho}\delta(\boldsymbol{\beta}; \boldsymbol{S}) + \mathcal{P}(\boldsymbol{\beta}). \tag{42}$$

By seeing the term $\frac{1}{\rho}\delta(\boldsymbol{\beta}; \boldsymbol{S})$ as the randomization term, the problem above can be seen as an example of a randomized procedure of the type we have considered so far.

Denoting the minimizer of the above equation as $\hat{\boldsymbol{\beta}} = \begin{pmatrix} \boldsymbol{\beta}_E \\ \mathbf{0} \end{pmatrix}$, the subgradient equation of the above becomes

$$\begin{aligned} \boldsymbol{\omega} := \frac{1}{\rho}\nabla\delta(\hat{\boldsymbol{\beta}}; \boldsymbol{S}) &= \nabla\ell(\hat{\boldsymbol{\beta}}; \boldsymbol{S}) + \partial\mathcal{P}(\hat{\boldsymbol{\beta}}) \\ &= \omega(\boldsymbol{D}, \boldsymbol{V}), \end{aligned}$$

where $\boldsymbol{D} = D(\boldsymbol{S})$ is a data dependent vector and $\boldsymbol{V} \in \mathcal{V}$ are naturally arising optimization variables restricted to a set $\mathcal{V}$. Choosing a model E is equivalent to having $\boldsymbol{V} \in \mathcal{V}$, hence to condition on the model it suffices to satisfy this constraint. Thus, to sample the data $\boldsymbol{D}$ post selection, i.e. conditional on (42) choosing model E, it suffices to sample $\boldsymbol{D}$ and $\boldsymbol{V}$ from the density as follows. Assuming that pre-selection $\boldsymbol{\omega} = \frac{1}{\rho}\nabla\delta(\hat{\boldsymbol{\beta}}; \boldsymbol{S})$ and $\boldsymbol{D}$ satisfy a CLT:

$$\begin{pmatrix} \boldsymbol{\omega} \\ \boldsymbol{D} \end{pmatrix} \xrightarrow{d} \mathcal{N}_{\dim(\boldsymbol{\omega})+\dim(\boldsymbol{D})}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$$

as $n \to \infty$, we sample $(\boldsymbol{D}, \boldsymbol{V})$ from the density proportional to standard multivariate Gaussian density

$$\phi_{\dim(\boldsymbol{\omega})+\dim(\boldsymbol{D})}\left(\boldsymbol{\Sigma}^{-1/2}\left(\begin{pmatrix} \omega(\boldsymbol{D}, \boldsymbol{V}) \\ \boldsymbol{D} \end{pmatrix} - \boldsymbol{\mu}\right)\right) \tag{43}$$

where the optimization variables are restricted to $\boldsymbol{V} \in \mathcal{V}$, assuming the Jacobian is a constant. In the examples later, we will choose the randomization $\boldsymbol{\omega}$ to be of mean zero and asymptotically independent of $\boldsymbol{D}$, in which case the density in (43) simplifies further as

$$\phi_{\dim(\boldsymbol{D})}\left(\boldsymbol{\Sigma}_{\boldsymbol{D}}^{-1/2}(\boldsymbol{D} - \boldsymbol{\mu}_{\boldsymbol{D}})\right) \cdot \phi_{\dim(\boldsymbol{\omega})}\left(\boldsymbol{\Sigma}_{\boldsymbol{\omega}}^{-1/2}\omega(\boldsymbol{D}, \boldsymbol{V})\right), \tag{44}$$

where $\boldsymbol{D} \xrightarrow{d} \mathcal{N}(\boldsymbol{\mu_D}, \boldsymbol{\Sigma_D})$ and $\boldsymbol{\omega} \xrightarrow{d} \mathcal{N}(\boldsymbol{0}, \boldsymbol{\Sigma_\omega})$ as $n \to \infty$. As before, if we are interested in parameter $\boldsymbol{\theta} = \theta(\mathbb{F})$ and have a target $\boldsymbol{T}$ satisfying the CLT pre-selection with mean $\boldsymbol{\theta}$, we decompose $\boldsymbol{D}$ in terms of $\boldsymbol{T}$ and the part of $\boldsymbol{D}$ independent of $\boldsymbol{T}$ pre-selection in the normal model. Similarly as before, we bootstrap $\boldsymbol{T}$ by sampling the bootstrap weights along with the optimization variables.

7.1. GLM example

To illustrate the idea in an example, let us do ℓ_1-penalized logistic regression. The data $\boldsymbol{S}$ above consists of the data matrix $\boldsymbol{X} \in \mathbb{R}^{n \times p}$ and a response vector $\boldsymbol{y} \in \mathbb{R}^n$ and the split data is $\boldsymbol{S}_1 = (\boldsymbol{X}_1, \boldsymbol{y}_1)$. The model selection objective on the split becomes

$$\underset{\boldsymbol{\beta}}{\text{minimize}}\, \ell(\boldsymbol{\beta}; (\boldsymbol{X}, \boldsymbol{y})) - \frac{1}{\rho}\delta(\boldsymbol{\beta}; (\boldsymbol{X}, \boldsymbol{y})) + \frac{\epsilon}{2}\|\boldsymbol{\beta}\|_2^2 + \lambda\|\boldsymbol{\beta}\|_1. \tag{45}$$

Since $\ell(\boldsymbol{\beta}; (\boldsymbol{X}, \boldsymbol{y})) - \frac{1}{\rho}\delta(\boldsymbol{\beta}; (\boldsymbol{X}, \boldsymbol{y})) = \frac{1}{\rho}\ell_1(\boldsymbol{\beta}; (\boldsymbol{X}_1, \boldsymbol{y}_1))$, the solution of the above optimization problem depends only on the split data. Denote the solution of (45) as $\hat{\boldsymbol{\beta}} = \begin{pmatrix} \boldsymbol{\beta}_E \\ \boldsymbol{0} \end{pmatrix}$. The KKT conditions of this problem are

$$\boldsymbol{\omega} := \frac{1}{\rho}\nabla\delta \begin{pmatrix} \boldsymbol{\beta}_E \\ \boldsymbol{0} \end{pmatrix} = \nabla\ell \begin{pmatrix} \boldsymbol{\beta}_E \\ \boldsymbol{0} \end{pmatrix} + \epsilon \begin{pmatrix} \boldsymbol{\beta}_E \\ \boldsymbol{0} \end{pmatrix} + \lambda \begin{pmatrix} \boldsymbol{s}_E \\ \boldsymbol{u}_{-E} \end{pmatrix} \tag{46}$$

with the constrains $\text{sign}(\boldsymbol{\beta}_E) = \boldsymbol{s}_E$ and $\|\boldsymbol{u}_{-E}\|_\infty \leqslant 1$. We want to sample the data vector (a function of $(\boldsymbol{X}, \boldsymbol{y})$) and the optimization variables $(\boldsymbol{\beta}_E, \boldsymbol{u}_{-E})$ such that they satisfy the KKT conditions corresponding to (45).

In order to get more explicit form for the randomization $\boldsymbol{\omega}$, we do the Taylor expansion of both gradient of δ and the gradient of the loss ℓ. Taylor expansion of the gradient of the loss around $\begin{pmatrix} \bar{\boldsymbol{\beta}}_E \\ \boldsymbol{0} \end{pmatrix}$, where $\bar{\boldsymbol{\beta}}_E$ is unpenalized MLE based on the predictors in $\boldsymbol{X}_E$, gives

$$\nabla\ell \begin{pmatrix} \boldsymbol{\beta}_E \\ \boldsymbol{0} \end{pmatrix} \approx \nabla\ell \begin{pmatrix} \bar{\boldsymbol{\beta}}_E \\ \boldsymbol{0} \end{pmatrix} + \nabla^2\ell \begin{pmatrix} \bar{\boldsymbol{\beta}}_E \\ \boldsymbol{0} \end{pmatrix} \begin{pmatrix} \boldsymbol{\beta}_E - \bar{\boldsymbol{\beta}}_E \\ \boldsymbol{0} \end{pmatrix}, \tag{47}$$

where $\bar{\boldsymbol{\beta}}_E$ solves the unpenalized problem with loss ℓ restricted to $\boldsymbol{X}_E \in \mathbb{R}^{n \times |E|}$. In the case of logistic loss, this becomes

$$\begin{aligned}
\nabla\ell \begin{pmatrix} \boldsymbol{\beta}_E \\ \boldsymbol{0} \end{pmatrix} &= -\boldsymbol{X}^T(\boldsymbol{y} - \pi(\boldsymbol{X}_E\boldsymbol{\beta}_E)) \\
&\approx -\boldsymbol{X}^T(\boldsymbol{y} - \pi(\boldsymbol{X}_E\bar{\boldsymbol{\beta}}_E)) + \boldsymbol{X}^T\boldsymbol{W}_E(\bar{\boldsymbol{\beta}}_E)\boldsymbol{X}_E(\boldsymbol{\beta}_E - \bar{\boldsymbol{\beta}}_E) \\
&= -\begin{pmatrix} \boldsymbol{X}_E^T\boldsymbol{W}_E(\bar{\boldsymbol{\beta}}_E)\boldsymbol{X}_E & \boldsymbol{0} \\ \boldsymbol{X}_{-E}^T\boldsymbol{W}_E(\bar{\boldsymbol{\beta}}_E)\boldsymbol{X}_E & \boldsymbol{I}_{p-|E|} \end{pmatrix} \boldsymbol{D} + \boldsymbol{X}^T\boldsymbol{W}_E(\bar{\boldsymbol{\beta}}_E)\boldsymbol{X}_E\boldsymbol{\beta}_E,
\end{aligned}$$

where $\boldsymbol{D}$ denotes

$$\boldsymbol{D} = \begin{pmatrix} \bar{\boldsymbol{\beta}}_E \\ \boldsymbol{X}_{-E}^T(\boldsymbol{y} - \pi(\boldsymbol{X}_E\bar{\boldsymbol{\beta}}_E)) \end{pmatrix}. \tag{48}$$

Taylor expansion of $\nabla\delta \begin{pmatrix} \boldsymbol{\beta}_E \\ \boldsymbol{0} \end{pmatrix}$ around the vector $\begin{pmatrix} \boldsymbol{\beta}_E^* \\ \boldsymbol{0} \end{pmatrix}$, where $\boldsymbol{\beta}_E^*$ are the population parameters corresponding the predictors in $\boldsymbol{X}_E$, gives

$$\nabla\delta \begin{pmatrix} \boldsymbol{\beta}_E \\ \boldsymbol{0} \end{pmatrix} \approx \nabla\delta \begin{pmatrix} \boldsymbol{\beta}_E^* \\ \boldsymbol{0} \end{pmatrix} + \nabla^2\delta \begin{pmatrix} \boldsymbol{\beta}_E^* \\ \boldsymbol{0} \end{pmatrix} \begin{pmatrix} \boldsymbol{\beta}_E - \boldsymbol{\beta}_E^* \\ \boldsymbol{0} \end{pmatrix}.$$

Since the Hessian above is approximately zero $\left(\nabla^2\delta\begin{pmatrix}\boldsymbol{\beta}_E^*\\ \mathbf{0}\end{pmatrix} = \rho\nabla^2\ell\begin{pmatrix}\boldsymbol{\beta}_E^*\\ \mathbf{0}\end{pmatrix} - \nabla^2\ell_1\begin{pmatrix}\boldsymbol{\beta}_E^*\\ \mathbf{0}\end{pmatrix} \approx 0\right)$, we have

$$\begin{aligned}\nabla\delta\begin{pmatrix}\boldsymbol{\beta}_E\\ \mathbf{0}\end{pmatrix} \approx \nabla\delta\begin{pmatrix}\boldsymbol{\beta}_E^*\\ \mathbf{0}\end{pmatrix} &= \rho\nabla\ell\begin{pmatrix}\boldsymbol{\beta}_E^*\\ \mathbf{0}\end{pmatrix} - \nabla\ell_1\begin{pmatrix}\boldsymbol{\beta}_E^*\\ \mathbf{0}\end{pmatrix}\\ &= -\rho\boldsymbol{X}^T(\boldsymbol{y} - \pi(\boldsymbol{X}_E\boldsymbol{\beta}_E^*)) + \boldsymbol{X}_1^T(\boldsymbol{y}_1 - \pi(\boldsymbol{X}_{1,E}\boldsymbol{\beta}_E^*)),\end{aligned}$$

implying

$$\boldsymbol{\omega} \approx \frac{1}{\rho}\nabla\delta\begin{pmatrix}\boldsymbol{\beta}_E^*\\ \mathbf{0}\end{pmatrix} = -\boldsymbol{X}^T(\boldsymbol{y} - \pi(\boldsymbol{X}_E\boldsymbol{\beta}_E^*)) + \frac{1}{\rho}\boldsymbol{X}_1^T(\boldsymbol{y}_1 - \pi(\boldsymbol{X}_{1,E}\boldsymbol{\beta}_E^*)). \tag{49}$$

A simple calculation shows the asymptotic covariance between $\boldsymbol{\omega}$ and $\boldsymbol{D}$ is zero, hence we sample from the density in (44), where all the covariances are estimated using bootstrap.

7.2. Multiple splits

Combining the ideas from multiple spits and data splitting, we now do inference after looking at the outcomes of the model selection procedures run on different splits of the data.

Let us in total do K splits of the data, each running a model selection procedure on the split data $(\boldsymbol{X}_k, \boldsymbol{y}_k) \in \mathbb{R}^{m_k\times p} \times \mathbb{R}^{m_k}$, of size m_k, selecting a model E_k at each of $k = 1, \ldots, K$. Let an analyst choose model E upon seeing the outcomes $E_1, \ldots, E_K$. Each of the model selection procedures can be written as

$$\ell(\boldsymbol{\beta}; (\boldsymbol{X}, \boldsymbol{y})) + \delta_k(\boldsymbol{\beta}; (\boldsymbol{X}, \boldsymbol{y})) + \mathcal{P}(\boldsymbol{\beta}), \tag{50}$$

where $\delta_k(\boldsymbol{\beta}; (\boldsymbol{X}; \boldsymbol{y})) = \ell(\boldsymbol{\beta}; (\boldsymbol{X}, \boldsymbol{y})) - \frac{1}{\rho_k}\ell_k(\boldsymbol{\beta}; \boldsymbol{X}_k, \boldsymbol{y}_k))$, $\rho_k = \frac{m_k}{n}$. Let us write the KKT conditions of (50) as

$$\boldsymbol{\omega}_k = \omega_k(\boldsymbol{D}, \boldsymbol{V}_k), \tag{51}$$

where $\boldsymbol{\omega}_k = \nabla\delta_k(\boldsymbol{\beta}; (\boldsymbol{X}_k, \boldsymbol{y}_k))$, $\boldsymbol{D}$ is the data vector and $\boldsymbol{V}_k \in \mathcal{V}_k$ are the optimization variables constrained to a set $\mathcal{V}_k$, $k = 1, \ldots, K$. Assuming that pre-selection the asymptotic covariance between any two of $\boldsymbol{D}, \boldsymbol{\omega}_1, \ldots, \boldsymbol{\omega}_K$, is zero and

$$\boldsymbol{D} - \boldsymbol{\mu}_{\boldsymbol{D}} \xrightarrow{d} \mathcal{N}_{\dim(\boldsymbol{D})}(\mathbf{0}, \boldsymbol{\Sigma}_{\boldsymbol{D}}), \quad \boldsymbol{\omega}_k \xrightarrow{d} \mathcal{N}_{\dim(\boldsymbol{\omega}_k)}(\mathbf{0}, \boldsymbol{\Sigma}_{\boldsymbol{\omega}_k}), \tag{52}$$

as $n \to \infty$, we have that the selective density of $\boldsymbol{D}, \boldsymbol{V}_1, \ldots, \boldsymbol{V}_K$, is

$$\phi_{\dim(\boldsymbol{D})}\left(\boldsymbol{\Sigma}_{\boldsymbol{D}}^{-1/2}(\boldsymbol{D} - \boldsymbol{\mu}_{\boldsymbol{D}})\right) \cdot \prod_{k=1}^{K}\phi_{\dim(\boldsymbol{\omega}_k)}\left(\boldsymbol{\Sigma}_{\boldsymbol{\omega}_k}^{-1/2}\omega_k(\boldsymbol{D}, \boldsymbol{V}_k)\right).$$

In the case of the ℓ_1-penalized logistic regression, the randomization reconstruction map becomes

$$\boldsymbol{\omega}_k = -\boldsymbol{X}^T(\boldsymbol{y} - \pi(\boldsymbol{X}_E\boldsymbol{\beta}_E^*)) + \frac{1}{\rho}\boldsymbol{X}_k^T(\boldsymbol{y}_k - \pi(\boldsymbol{X}_{k,E}\boldsymbol{\beta}_E^*)), \tag{53}$$

for $k = 1, \ldots, K$ and in this example the cross covariances between $\boldsymbol{D}, \boldsymbol{\omega}_1, \ldots, \boldsymbol{\omega}_K$ are zero. As elaborated in detail in multiple views of the data for GLMs (Section 6.2), we do inference in the same way. The implementation results of this procedure are presented in Section 8.1.

8. Sampling via projected Langevin MCMC

To sample from either the plugin CLT or the bootstrap density, we use projected Langevin MC to sample from a log-concave density with constraints. Given a convex set $K \subset \mathbb{R}^n$ with a nonempty

interior, let us denote with $\mathcal{P}_K$ the projection onto K. We are interested in sampling from a density $f(\boldsymbol{x})$ on $\mathbb{R}^n$ given by

$$\frac{df(\boldsymbol{x})}{d\boldsymbol{x}} \propto e^{-\tilde{f}(\boldsymbol{x})} \mathbb{I}_{\{\boldsymbol{x}\in K\}},$$

where $\tilde{f} : K \to \mathbb{R}$ is a convex and differentiable function (or at least sub-differentiable).

Based on the previous point $\boldsymbol{X}_k$ in the chain, projected Langevin computes the next point using the following update

$$\boldsymbol{X}_{k+1} = \mathcal{P}_K \left(\boldsymbol{X}_k - \eta \nabla \tilde{f}(\boldsymbol{X}_k) + \sqrt{2\eta}\boldsymbol{\xi}_k \right),$$

where $\boldsymbol{\xi}_1, \boldsymbol{\xi}_2, \ldots$ are i.i.d. sequence of standard normal in $\mathbb{R}^n$ and η is the step size. Since our constraint region is simple due to the change of variables of [Tian Harris et al., 2016], the projection step is easy. [Bubeck et al., 2015] proves that this chain will converge to a true density under some conditions.

Taking the randomization density g and the bootstrap weights density f_α to be log-concave, we use the updates above to sample from the bootstrap densities. Let us now write the Langevin updates for the sampling of density some of the examples mentioned so far. We write the updates for both plugin CLT and the bootstrap sampling densities.

Example 8.1 (Lasso, Section 5.1) *Let us first write down the projected Langevin update to sample from the plugin CLT density in (16)*

$$\begin{pmatrix} \boldsymbol{T}^{(i+1)} \\ \boldsymbol{\beta}_E^{(i+1)} \\ \boldsymbol{u}_{-E}^{(i+1)} \end{pmatrix} = \begin{pmatrix} \boldsymbol{T}^{(i)} - \eta \widehat{\boldsymbol{\Sigma}}_{\boldsymbol{T}}^{-1}(\boldsymbol{T}^{(i)} - \boldsymbol{\theta}) + \eta \widetilde{\boldsymbol{M}}^T \nabla_{\boldsymbol{\omega}} \log g(\omega(\boldsymbol{T}^{(i)}, \boldsymbol{\beta}_E^{(i)}, \boldsymbol{u}_{-E}^{(i)})) + \sqrt{2\eta}\boldsymbol{\xi}^{(i)} \\ \mathcal{P}_1 \left(\boldsymbol{\beta}_E^{(i)} + \eta \boldsymbol{B}^T \nabla_{\boldsymbol{\omega}} \log g(\omega(\boldsymbol{T}^{(i)}, \boldsymbol{\beta}_E^{(i)}, \boldsymbol{u}_{-E}^{(i)})) + \sqrt{2\eta}\boldsymbol{\xi}_1^{(i)} \right) \\ \mathcal{P}_2 \left(\boldsymbol{u}_{-E}^{(i)} + \eta \boldsymbol{U}^T \nabla_{\boldsymbol{\omega}} \log g(\omega(\boldsymbol{T}^{(i)}, \boldsymbol{\beta}_E^{(i)}, \boldsymbol{u}_{-E}^{(i)})) + \sqrt{2\eta}\boldsymbol{\xi}_2^{(i)} \right) \end{pmatrix},$$

where $\omega(\boldsymbol{T}, \boldsymbol{\beta}_E, \boldsymbol{u}_{-E}) = \widetilde{\boldsymbol{M}}\boldsymbol{T} + \boldsymbol{B}\boldsymbol{\beta}_E + \boldsymbol{U}\boldsymbol{u}_{-E} + \widetilde{\boldsymbol{L}}$ *and* $(\boldsymbol{\xi}^{(i)}, \boldsymbol{\xi}_1^{(i)}, \boldsymbol{\xi}_2^{(i)}) \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{I}_a) \times \mathcal{N}(\boldsymbol{0}, \boldsymbol{I}_{|E|}) \times \mathcal{N}(\boldsymbol{0}, \boldsymbol{I}_{p-|E|})$ *and independent of all the other random variables. Here,* $\mathcal{P}_1$ *is the projection onto the orthant* $\mathbb{R}^{|E|}_{\boldsymbol{s}_E}$ *and* $\mathcal{P}_2$ *is the projection onto the cube* $[-1, 1]^{p-|E|}$

Now let us write down the projected Langevin update to sample from the bootstrap density in (18). Based on the current point $(\boldsymbol{\alpha}^{(i)}, \boldsymbol{\beta}_E^{(i)}, \boldsymbol{u}_{-E}^{(i)})$*, the update becomes*

$$\begin{pmatrix} \boldsymbol{\alpha}^{(i+1)} \\ \boldsymbol{\beta}_E^{(i+1)} \\ \boldsymbol{u}_{-E}^{(i+1)} \end{pmatrix} = \begin{pmatrix} \mathcal{P} \left(\boldsymbol{\alpha}^{(i)} + \eta \nabla_{\boldsymbol{\alpha}} \left(\sum_{j=1}^n \log f_\alpha(\alpha_j^{(i)}) \right) + \eta \overline{\boldsymbol{M}}^T \nabla_{\boldsymbol{\omega}} \log g \left(\omega^b(\boldsymbol{\alpha}^{(i)}, \boldsymbol{\beta}_E^{(i)}, \boldsymbol{u}_{-E}^{(i)}) \right) + \sqrt{2\eta}\boldsymbol{\xi}^{(i)} \right) \\ \mathcal{P}_1 \left(\boldsymbol{\beta}_E^{(i)} + \eta \boldsymbol{B}^T \nabla_{\boldsymbol{\omega}} \log g(\omega^b(\boldsymbol{\alpha}^{(i)}, \boldsymbol{\beta}_E^{(i)}, \boldsymbol{u}_{-E}^{(i)})) + \sqrt{2\eta}\boldsymbol{\xi}_1^{(i)} \right) \\ \mathcal{P}_2 \left(\boldsymbol{u}_{-E}^{(i)} + \eta \boldsymbol{U}^T \nabla_{\boldsymbol{\omega}} \log g(\omega^b(\boldsymbol{\alpha}^{(i)}, \boldsymbol{\beta}_E^{(i)}, \boldsymbol{u}_{-E}^{(i)})) + \sqrt{2\eta}\boldsymbol{\xi}_2^{(i)} \right) \end{pmatrix},$$

where $\omega^b(\boldsymbol{\alpha}, \boldsymbol{\beta}_E, \boldsymbol{u}_{-E}) = \overline{\boldsymbol{M}}\boldsymbol{\alpha} + \boldsymbol{B}\boldsymbol{\beta}_E + \boldsymbol{U}\boldsymbol{\beta}u_{-E} + \widetilde{\boldsymbol{L}} + \widetilde{\boldsymbol{M}}\boldsymbol{\theta}$*,* $\overline{\boldsymbol{M}} = \widetilde{\boldsymbol{M}}\boldsymbol{A}^T(\boldsymbol{X}_E^T\boldsymbol{X}_E)^{-1}\boldsymbol{X}_E^T\widehat{\boldsymbol{\varepsilon}}$*, given in (17) and* $(\boldsymbol{\xi}^{(i)}, \boldsymbol{\xi}_1^{(i)}, \boldsymbol{\xi}_2^{(i)}) \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{I}_n) \times \mathcal{N}(\boldsymbol{0}, \boldsymbol{I}_{|E|}) \times \mathcal{N}(\boldsymbol{0}, \boldsymbol{I}_{p-|E|})$ *and independent of all the other random variables. Here,* $\mathcal{P}$ *is the projection onto* $(\mathrm{supp}(\mathbb{F}_\alpha))^n$*.*

Example 8.2 (Forward stepwise, Section 6.1) *To sample from the plugin CLT density in (29), the Langevin update at* $(i+1)$*-th step based on* $(\boldsymbol{T}^{(i)}, \boldsymbol{z}_1^{(i)}, \ldots, \boldsymbol{z}_K^{(i)})$*, the point at the* i*-th step, is the following*

$$\begin{pmatrix} \boldsymbol{T}^{(i+1)} \\ \boldsymbol{z}_1^{(i+1)} \\ \vdots \\ \boldsymbol{z}_K^{(i+1)} \end{pmatrix} = \begin{pmatrix} \boldsymbol{T}^{(i)} - \eta \widehat{\boldsymbol{\Sigma}}_{\boldsymbol{T}}^{-1}(\boldsymbol{T}^{(i)} - \boldsymbol{\theta}) + \eta \sum_{k=1}^K \widetilde{\boldsymbol{M}}_k^T \nabla_{\boldsymbol{\omega}_k} \log g(\omega_k(\boldsymbol{T}^{(i)}, \boldsymbol{z}_k^{(i)})) + \sqrt{2\eta}\boldsymbol{\xi}^{(i)} \\ \mathcal{P}_1 \left(\boldsymbol{z}_1^{(i)} + \eta \nabla_{\boldsymbol{\omega}_1} \log g_1(\omega_1(\boldsymbol{T}^{(i)}, \boldsymbol{z}_1^{(i)})) + \sqrt{2\eta}\boldsymbol{\xi}_1^{(i)} \right) \\ \vdots \\ \mathcal{P}_K \left(\boldsymbol{z}_K^{(i)} + \eta \nabla_{\boldsymbol{\omega}_K} \log g_K(\omega_K(\boldsymbol{T}^{(i)}, \boldsymbol{z}_K^{(i)}) + \sqrt{2\eta}\boldsymbol{\xi}_K^{(i)} \right) \end{pmatrix},$$

where $\omega_k(\boldsymbol{z}_k, \boldsymbol{T}) = \boldsymbol{z}_k + \boldsymbol{M}_k\boldsymbol{F} + \widetilde{\boldsymbol{M}}_k\boldsymbol{T}$ *and* $(\boldsymbol{\xi}^{(i)}, \boldsymbol{\xi}_1^{(i)}, \ldots, \boldsymbol{\xi}_K^{(i)}) \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{I}_a) \times \mathcal{N}(\boldsymbol{0}, \boldsymbol{I}_p) \times \ldots \times \mathcal{N}(\boldsymbol{0}, \boldsymbol{I}_{p-K+1})$ *and independent of all the other variables. Here,* $\mathcal{P}_k$ *denotes the projection onto* $\partial I_{\mathcal{B}_k}(\hat{\boldsymbol{\eta}}_k)$ *for* $k = 1, \ldots, K$.

To sample from the bootstrap density in (31)*, the Langevin update at* $(i+1)$*-th step based on* $(\boldsymbol{\alpha}^{(i)}, \boldsymbol{z}_1^{(i)}, \ldots, \boldsymbol{z}_K^{(i)})$*, the point at the i-th step, is the following*

$$\begin{pmatrix} \boldsymbol{\alpha}^{(i+1)} \\ \boldsymbol{z}_1^{(i+1)} \\ \vdots \\ \boldsymbol{z}_K^{(i+1)} \end{pmatrix} = \begin{pmatrix} \mathcal{P}\left(\boldsymbol{\alpha}^{(i)} + \eta\nabla_{\boldsymbol{\alpha}}\left(\sum_{j=1}^n \log f(\alpha_j^{(i)})\right) + \eta\sum_{k=1}^K \overline{\boldsymbol{M}}_k^T\nabla_{\boldsymbol{\omega}_k}\log g(\omega_k^b(\boldsymbol{\alpha}^{(i)}, \boldsymbol{z}_k^{(i)})) + \sqrt{2\eta}\boldsymbol{\xi}^{(i)}\right) \\ \mathcal{P}_1\left(\boldsymbol{z}_1^{(i)} + \eta\nabla_{\boldsymbol{\omega}_1}\log g_1(\omega_1^b(\boldsymbol{\alpha}^{(i)}, \boldsymbol{z}_1^{(i)})) + \sqrt{2\eta}\boldsymbol{\xi}_1^{(i)}\right) \\ \vdots \\ \mathcal{P}_K\left(\boldsymbol{z}_K^{(i)} + \eta\nabla_{\boldsymbol{\omega}_K}\log g_K(\omega_K^b(\boldsymbol{\alpha}^{(i)}, \boldsymbol{z}_K^{(i)}) + \sqrt{2\eta}\boldsymbol{\xi}_K^{(i)}\right) \end{pmatrix},$$

where ω_k^b *is given in* (30) *and* $(\boldsymbol{\xi}^{(i)}, \boldsymbol{\xi}_1^{(i)}, \ldots, \boldsymbol{\xi}_K^{(i)}) \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{I}_n) \times \mathcal{N}(\boldsymbol{0}, \boldsymbol{I}_p) \times \ldots \times \mathcal{N}(\boldsymbol{0}, \boldsymbol{I}_{p-K+1})$ *and independent of all the other variables. Here,* $\mathcal{P}$ *denotes the projection onto* $(\mathrm{supp}(\mathbb{F}_\alpha))^n$.

Example 8.3 (Multiple views with GLMs and group Lasso penalty 6.2) *Now we sample from the density in* (37) *using Langevin samples where at each step we move the target vector* $\boldsymbol{T}$ *and the optimization variables* $\{(\gamma_{k,g})_{g\in A_k}, (\boldsymbol{z}_{k,h})_{h\in -A_k}\}_{k=1}^K$.

The gradient of the logarithm of the selective density in(37) *becomes*

$$G\begin{pmatrix} \boldsymbol{T} \\ (\gamma_{1,g})_{g\in A_1} \\ (\boldsymbol{z}_{1,h})_{h\in -A_1} \\ \vdots \\ (\gamma_{K,g})_{g\in A_K} \\ (\boldsymbol{z}_{K,h})_{h\in -A_K} \end{pmatrix} = \begin{pmatrix} -\widehat{\boldsymbol{\Sigma}}_{\boldsymbol{T}}^{-1}(\boldsymbol{T} - \boldsymbol{\theta}) + \sum_{k=1}^K \left(\widetilde{\boldsymbol{M}}_k^T\nabla\log g_k(\boldsymbol{\omega}_k) + \nabla_{\boldsymbol{T}}\log J_k\right) \\ \boldsymbol{\Gamma}_1^T\nabla\log g_1(\boldsymbol{\omega}_1) + \nabla_{(\gamma_{1,g})_{g\in A_1}}\log J_1 \\ \boldsymbol{Z}_1^T\nabla\log g_1(\boldsymbol{\omega}_1) + \nabla_{(\boldsymbol{z}_{1,h})_{h\in -A_1}}\log J_1 \\ \vdots \\ \boldsymbol{\Gamma}_K^T\nabla\log g_K(\boldsymbol{\omega}_K) + \nabla_{(\gamma_{K,g})_{g\in A_K}}\log J_K \\ \boldsymbol{Z}_K^T\nabla\log g_K(\boldsymbol{\omega}_K) + \nabla_{(\boldsymbol{z}_{K,h})_{h\in -A_K}}\log J_K \end{pmatrix},$$

where $\boldsymbol{\omega}_k$ *is the randomization reconstruction from* (36)*,* $k = 1, \ldots, K$.

The Langevin update at step i becomes:

$$\begin{pmatrix} \boldsymbol{T} \\ (\gamma_{1,g})_{g\in A_1} \\ (\boldsymbol{z}_{1,h})_{h\in -A_1} \\ \vdots \\ (\gamma_{K,g})_{g\in A_K} \\ (\boldsymbol{z}_{K,h})_{h\in -A_K} \end{pmatrix}^{(i+1)} = \mathcal{P}\left(\begin{pmatrix} \boldsymbol{T} \\ (\gamma_{1,g})_{g\in A_1} \\ (\boldsymbol{z}_{1,h})_{h\in -A_1} \\ \vdots \\ (\gamma_{K,g})_{g\in A_K} \\ (\boldsymbol{z}_{K,h})_{h\in -A_K} \end{pmatrix}^{(i)} + \frac{\eta}{2}G\left(\begin{pmatrix} \boldsymbol{T} \\ (\gamma_{1,g})_{g\in A_1} \\ (\boldsymbol{z}_{1,h})_{h\in -A_1} \\ \vdots \\ (\gamma_{K,g})_{g\in A_K} \\ (\boldsymbol{z}_{K,h})_{h\in -A_K} \end{pmatrix}^{(i)}\right) + \sqrt{\eta}\boldsymbol{\xi}_i\right), \tag{54}$$

where η *is the step size and* $\boldsymbol{\xi}_i \overset{i.i.d.}{\sim} \mathcal{N}_{a+pK}(0, I_{a+pK})$*, independent of all other variables. Here* $\mathcal{P}$ *is a projection which projects each* $\gamma_{k,g}$ *coordinate,* $g \in A_k$*, onto positive axis and each* $\boldsymbol{z}_{k,h}$ *coordinate,* $h \in -A_k$*, onto a ball in* $\mathbb{R}^{|h|}$ *centered at the origin and of radius* $\lambda_{k,g}$ *for all* $k = 1, \ldots, K$.

To sample from the bootstrap density in (39) *using Langevin samples where at each step we move the bootstrap weights, vector* $\boldsymbol{\alpha} \in \mathbb{R}^n$*, and the optimization variables* $\{(\gamma_{k,g})_{g\in A_k}, (\boldsymbol{z}_{k,h})_{h\in -A_k}\}_{k=1}^K$. *The gradient of the logarithm of the selective density from* (37) *becomes*

$$G^b\begin{pmatrix} \boldsymbol{\alpha} \\ (\gamma_{1,g})_{g\in A_1} \\ (\boldsymbol{z}_{1,h})_{h\in -A_1} \\ \vdots \\ (\gamma_{K,g})_{g\in A_K} \\ (\boldsymbol{z}_{K,h})_{h\in -A_K} \end{pmatrix} = \begin{pmatrix} (\nabla_{\boldsymbol{\alpha}}\sum_{i=1}^n f_\alpha(\alpha_i)) + \sum_{k=1}^K \left(\overline{\boldsymbol{M}}_k^T\nabla\log g_k(\boldsymbol{\omega}_k) + \nabla_{\boldsymbol{T}}\log J_k\right) \\ \boldsymbol{\Gamma}_1^T\nabla\log g_1(\boldsymbol{\omega}_1) + \nabla_{(\gamma_{1,g})_{g\in A_1}}\log J_1 \\ \boldsymbol{Z}_1^T\nabla\log g_1(\boldsymbol{\omega}_1) + \nabla_{(\boldsymbol{z}_{1,h})_{h\in -A_1}}\log J_1 \\ \vdots \\ \boldsymbol{\Gamma}_K^T\nabla\log g_K(\boldsymbol{\omega}_K) + \nabla_{(\gamma_{K,g})_{g\in A_K}}\log J_K \\ \boldsymbol{Z}_K^T\nabla\log g_K(\boldsymbol{\omega}_K) + \nabla_{(\boldsymbol{z}_{K,h})_{h\in -A_K}}\log J_K \end{pmatrix},$$

where $\boldsymbol{\omega}_k$ is the randomization reconstruction from (36), $k = 1, \ldots, K$. In the case of standard normal bootstrap weights $\nabla_{\boldsymbol{\alpha}} \sum_{i=1}^n f_\alpha(\alpha_i) = -\boldsymbol{\alpha}$.

The Langevin update is analogous to (54) with the addition of projection of the bootstrap weights $\boldsymbol{\alpha}$ onto their distribution support, $\text{supp}(\mathbb{F}_\alpha)^n$, in the case when the bootstrap weights are not of full support.

Remark 12 *There are both theoretical and practical considerations when choosing the density of the randomization, g, and the density of bootstrap weights, f_α. We take g to be Laplace, logistic or Gaussian and f_α to be standard normal.*

8.1. Implementation results

We present the empirical results using the proposed bootstrap and the plugin CLT test statistics in the ℓ_1-penalized logistic model, where each group is individual variable. At each iteration a new $\boldsymbol{X} \in \mathbb{R}^{n \times p}$, $\boldsymbol{y} \in \mathbb{R}^n$ instance is created from the logistic model with y_i taking values 0 and 1 with probabilities $1 - \pi(\boldsymbol{x}_i^T \boldsymbol{\beta})$ and $\pi(\boldsymbol{x}_i^T \boldsymbol{\beta})$, respectively, for $i = 1, \ldots, n$. The true underlying parameter vector $\boldsymbol{\beta} \in \mathbb{R}^p$ is taken so that its true sparsity is s. For each $(\boldsymbol{X}, \boldsymbol{y})$ instance, we select model E by adding the randomization to the ℓ_1-penalized objective and construct the confidence intervals for each coefficient $\beta^*_{j|E}$, $j = 1, \ldots, |E|$, in the selected model by using both the plugin CLT and the bootstrap sampler. [21] The coverage is computed by averaging over many iterations (instances). Also, we construct the naive confidence intervals based on the asymptotic normality of $\bar{\boldsymbol{\beta}}_E$ that do no take the selection into account.

The implemented examples include:

- One view of the data after adding Laplace randomization (Figure 3) or Gaussian randomization (Figure 4) to the ℓ_1-penalized objective function.
- Multiple views of the data (Figure 5). We do three model selection procedures (views). At each view we keep the same loss and penalty while drawing new independent randomization each time. The selected set E is chosen to be the union of the models selected at each step. We construct the confidence intervals for each of the coefficients in the model E.
- Data carving vs data splitting (Figure 6a). Here we do one split of the data keeping 80% of the data for model selection and construct both classical data splitting intervals and the data carving intervals described here (Section 7). Standard data splitting intervals are constructed using only leftover 20% of the data while data carving (plugin CLT and bootstrap intervals) uses the whole data set conditional on the model selected based on the first part of the data. We see that the data splitting intervals are 1.42 times larger than the intervals constructed using either CLT plugin pivot or bootstrap pivot.
- Multiple splits of the data (Figure 6b). Here we do three splits of the data, selecting the model on the random 50% of the data each time. The inference is then done for all the coefficients corresponding to the union of all three selected models. After three random splits, there would with high probability be no leftover data hence the data splitting approach would not be doable here.

All the code for the examples mentioned and more, including forward stepwise and marginal screening, is available at `https://github.com/jonathan-taylor/selective-inference`.

[21] An important thing is that we run a sampler (either plugin CLT or bootstrap) by taking the reference parameter to be unpenalized MLE ($\bar{\boldsymbol{\beta}}_E$) and get the samples at other values of the parameter $\boldsymbol{\beta}$, including the truth $\boldsymbol{\beta}^*_E$, by tilting the samples at the original reference. This way for each $(\boldsymbol{X}, \boldsymbol{y})$ instance we run the Monte Carlo chain only once and get the pivots at all parameter values. Then the confidence intervals are constructed by inverting these tests since the pivots are uniformly distributed under the null.

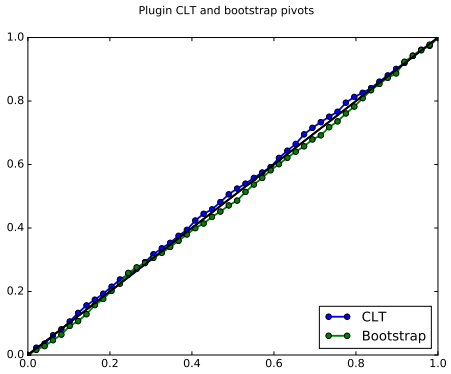

(a) $n = 300, p = 20, s = 3$, CLT coverage: 0.90, Boot: 0.90, Naive: 0.87

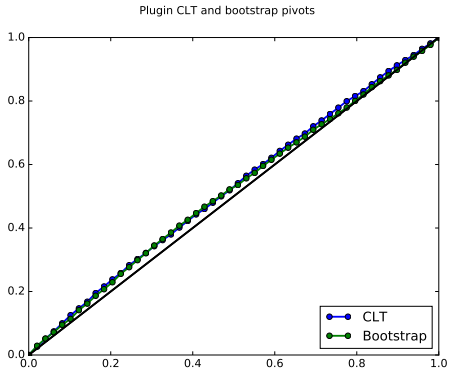

(b) $n = 300, p = 100, s = 5$, CLT coverage: 0.88, Boot: 0.89, Naive: 0.84

Figure 3: The pivots constructed for testing the true coefficients in the selected model ($H_0 : \beta^*_{j|E}$, $j = 1, \dots |E|$) after adding Laplace randomization (i.i.d. Laplace with mean 0 and scale 1) to the objective.

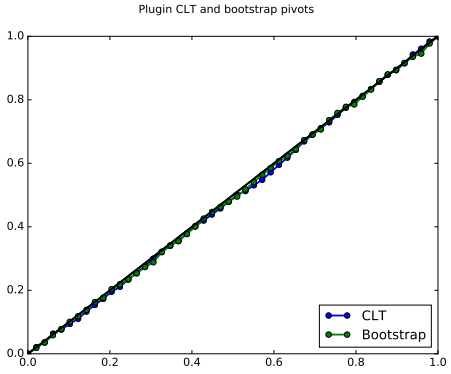

(a) $n = 300, p = 20, s = 3$, CLT coverage: 0.90, Boot: 0.90, Naive: 0.87

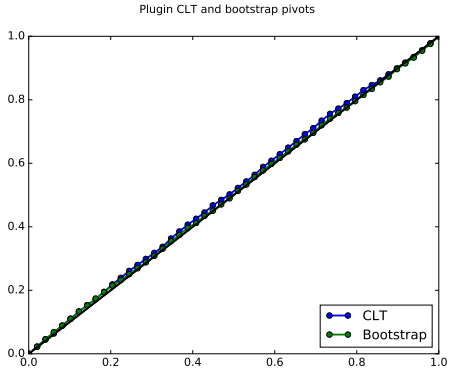

(b) $n = 300, p = 100, s = 5$, CLT coverage: 0.89, Boot: 0.89, Naive: 0.83

Figure 4: The pivots constructed for testing the true coefficients in the selected model ($H_0 : \beta^*_{j|E}$, $j = 1, \dots |E|$) after adding Gaussian randomization (i.i.d. standard Gaussians) to the objective.

Acknowledgements

Jelena Markovic was supported by Stanford Graduate Fellowship. Jonathan Taylor was supported in part by National Science Foundation grant DMS-1208857 and Air Force Office of Sponsored Research grant 113039.

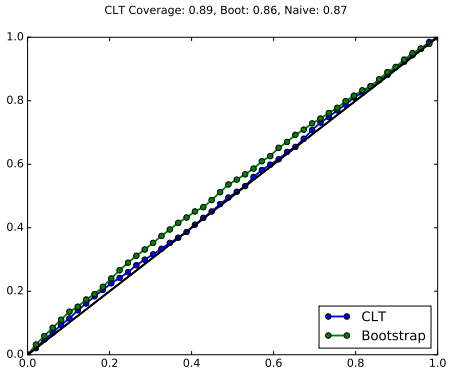

Figure 5: Multiple views of the data. The graph represents the p-values constructed for testing each of the coefficients in the union of the three selected models. Each model is selected by adding Laplace randomization to the objective ($n = 300, p = 20, s = 0$).

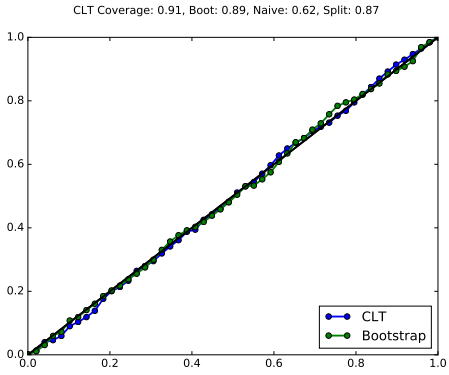

(a) Data spliting vs data carving

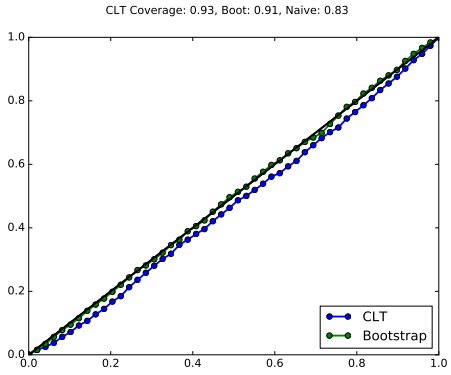

(b) Multiple splits

Figure 6: (a) After one split of the data, the intervals are constructed for all the coefficients in the selected model. The plot represents the p-values for testing the coefficients at zero. (b) After three splits of the data, the intervals are constructed for all the coefficients in the union of the three selected models. The plot represents the p-values for testing the coefficients are zero. $n = 300, p = 20, s = 0$ in both examples.

References

[Berk et al., 2013] Berk, R., Brown, L., Buja, A., Zhang, K., and Zhao, L. (2013). Valid post-selection inference. *The Annals of Statistics*, 41(2):802–837.

[Bubeck et al., 2015] Bubeck, S., Eldan, R., and Lehec, J. (2015). Sampling from a log-concave distribution with Projected Langevin Monte Carlo. *arXiv preprint arXiv:1507.02564*.

[Buja et al., 2014] Buja, A., Berk, R., Brown, L., George, E., Pitkin, E., Traskin, M., Zhang, K., Zhao, L., and White, D. T. H. (2014). A conspiracy of random predictors and model violations against classical inference in regression. *arXiv preprint*.

[Chatterjee, 2005] Chatterjee, S. (2005). A simple invariance theorem. *arXiv preprint math/0508213*.

[Efron, 1979] Efron, B. (1979). Bootstrap methods: another look at the jackknife. *The Annals of Statistics*, 7(1):1–26.

[Fithian et al., 2014] Fithian, W., Sun, D., and Taylor, J. (2014). Optimal inference after model selection. *arXiv preprint arXiv:1410.2597*.

[Freedman et al., 1981] Freedman, D. A. et al. (1981). Bootstrapping regression models. *The Annals of Statistics*, 9(6):1218–1228.

[Hanley and MacGibbon, 2006] Hanley, J. A. and MacGibbon, B. (2006). Creating non-parametric bootstrap samples using poisson frequencies. *computer methods and programs in biomedicine*, 83(1):57–62.

[Hansen, 1999] Hansen, B. E. (1999). The grid bootstrap and the autoregressive model. *Review of Economics and Statistics*, 81(4):594–607.

[Laber et al., 2014] Laber, E. B., Lizotte, D. J., Qian, M., Pelham, W. E., and Murphy, S. A. (2014). Dynamic treatment regimes: Technical challenges and applications. *Electronic journal of statistics*, 8(1):1225.

[Lee et al., 2016] Lee, J. D., Sun, D. L., Sun, Y., and Taylor, J. E. (2016). Exact post-selection inference, with application to the lasso. *The Annals of Statistics*, 44(3):907–927.

[Lee and Taylor, 2014] Lee, J. D. and Taylor, J. E. (2014). Exact post model selection inference for marginal screening. In *Advances in Neural Information Processing Systems*, pages 136–144.

[Leeb, 2015] Leeb, H. (2015). Comment. *Journal of the American Statistical Association*, 110(512):1457–1459.

[Leeb and Pötscher, 2006a] Leeb, H. and Pötscher, B. M. (2006a). Can one estimate the conditional distribution of post-model-selection estimators? *The Annals of Statistics*, 34(5):2554–2591.

[Leeb and Pötscher, 2006b] Leeb, H. and Pötscher, B. M. (2006b). Performance limits for estimators of the risk or distribution of shrinkage-type estimators, and some general lower risk-bound results. *Econometric Theory*, 22(01):69–97.

[Liu, 1988] Liu, R. Y. (1988). Bootstrap procedures under some non-iid models. *The Annals of Statistics*, 16(4):1696–1708.

[Loftus and Taylor, 2015] Loftus, J. R. and Taylor, J. E. (2015). Selective inference in regression models with groups of variables. *arXiv preprint arXiv:1511.01478*.

[Mammen, 1993] Mammen, E. (1993). Bootstrap and wild bootstrap for high dimensional linear models. *The Annals of Statistics*, 21(1):255–285.

[McKeague and Qian, 2015] McKeague, I. W. and Qian, M. (2015). An adaptive resampling test for detecting the presence of significant predictors. *Journal of the American Statistical Association*, 110(512):1422–1433.

[Praestgaard and Wellner, 1993] Praestgaard, J. and Wellner, J. A. (1993). Exchangeably weighted bootstraps of the general empirical process. *The Annals of Probability*, 21(4):2053–2086.

[Romano et al., 2012] Romano, J. P., Shaikh, A. M., et al. (2012). On the uniform asymptotic validity of subsampling and the bootstrap. *The Annals of Statistics*, 40(6):2798–2822.

[Rubin et al., 1981] Rubin, D. B. et al. (1981). The Bayesian bootstrap. *The Annals of Statistics*, 9(1):130–134.

[Taylor et al., 2013] Taylor, J., Loftus, J., and Tibshirani, R. (2013). Tests in adaptive regression via the Kac-Rice formula. *The Annals of Statistics*, 44(2):743–770.

[Tian et al., 2016] Tian, X., Bi, N., and Taylor, J. (2016). Magic: a general, powerful and tractable method for selective inference. *arXiv preprint arXiv:1607.02630*.

[Tian and Taylor, 2015a] Tian, X. and Taylor, J. (2015a). Asymptotics of selective inference. *arXiv preprint arXiv:1501.03588*.

[Tian and Taylor, 2015b] Tian, X. and Taylor, J. E. (2015b). Selective inference with a randomized response. *arXiv preprint arXiv:1507.06739*.

[Tian Harris et al., 2016] Tian Harris, X., Snigdha, P., Markovic, J., Bi, N., and Taylor, J. (2016). Selective sampling after solving a convex problem. *arXiv preprint arXiv:1609.05609*.

[Tibshirani et al., 2015] Tibshirani, R. J., Rinaldo, A., Tibshirani, R., and Wasserman, L. (2015). Uniform asymptotic inference and the bootstrap after model selection. *arXiv preprint arXiv:1506.06266*.

[Wu, 1986] Wu, C.-F. J. (1986). Jackknife, bootstrap and other resampling methods in regression analysis. *The Annals of Statistics*, 14(4):1261–1295.

[Zou and Hastie, 2005] Zou, H. and Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B*, 67(2):301–320.

Appendices

A. Selective CLT - version 1

Here we prove a version of selective CLT that does not assume a asymptotic linearity condition on the statistic $\boldsymbol{T}$; however, it requires the unknown parameter and the selection region to converge as the sample size n increase, hence not accounting for the rare selection events.

Denote the pivot as

$$\begin{aligned}
&\mathcal{P}\left(\sqrt{n}(\boldsymbol{T}-\boldsymbol{\mu}), \boldsymbol{A}^T\boldsymbol{\mu}, \sqrt{n}\boldsymbol{A}_M\boldsymbol{\mu}, \boldsymbol{A}_M, \boldsymbol{b}_M\right)\\
&= (\mathbb{P}_{\boldsymbol{G}}\times\mathbb{G})\big\{\|\boldsymbol{\Sigma}_{\boldsymbol{A}}^{-1/2}\boldsymbol{G}\|_2 \geqslant \|\sqrt{n}\boldsymbol{\Sigma}_{\boldsymbol{A}}^{-1/2}\boldsymbol{A}^T(\boldsymbol{T}-\boldsymbol{\mu})\|_2 \,\big|\\
&\qquad\qquad \sqrt{n}\boldsymbol{A}_M\boldsymbol{T}_{\boldsymbol{A}} + \boldsymbol{A}_M\boldsymbol{C}(\boldsymbol{G}+\sqrt{n}\boldsymbol{A}^T\boldsymbol{\mu}) + \boldsymbol{\omega} \geqslant \boldsymbol{b}_M, \boldsymbol{T}_{\boldsymbol{A}}\big\},\\
&= (\mathbb{P}_{\boldsymbol{G}}\times\mathbb{G})\big\{\|\boldsymbol{\Sigma}_{\boldsymbol{A}}^{-1/2}\boldsymbol{G}\|_2 \geqslant \|\sqrt{n}\boldsymbol{\Sigma}_{\boldsymbol{A}}^{-1/2}\boldsymbol{A}^T(\boldsymbol{T}-\boldsymbol{\mu})\|_2 \,\big|\\
&\qquad\qquad \sqrt{n}\boldsymbol{A}_M(\boldsymbol{T}_{\boldsymbol{A}}-\boldsymbol{\mu}_{\boldsymbol{A}}) + \boldsymbol{A}_M\boldsymbol{C}\boldsymbol{G} + \sqrt{n}\boldsymbol{A}_M\boldsymbol{\mu} + \boldsymbol{\omega} \geqslant \boldsymbol{b}_M, \boldsymbol{T}_{\boldsymbol{A}}\big\},
\end{aligned}$$

where $\mathbb{P}_{\boldsymbol{G}}$ is under $\boldsymbol{G}\sim\mathcal{N}_a(\boldsymbol{0},\boldsymbol{\Sigma}_{\boldsymbol{A}})$. Here $\boldsymbol{A}_M\in\mathbb{R}^{d\times p}$ and $\boldsymbol{b}_M\in\mathbb{R}^d$.[22]

To prove a version of selective CLT, let us introduce a new set of assumptions.

- Convergence of the selection regions:

$$\lim_{n\to\infty}\boldsymbol{A}_M = \widetilde{\boldsymbol{A}} \quad\text{and}\quad \lim_{n\to\infty}\boldsymbol{b}_M = \tilde{\boldsymbol{b}} \tag{55}$$

  for some $\widetilde{\boldsymbol{A}}\in\mathbb{R}^{d\times p}$ and $\tilde{\boldsymbol{b}}\in\mathbb{R}^d$.
- Convergence of the parameter:

$$\lim_{n\to\infty}\sqrt{n}\boldsymbol{A}_M\boldsymbol{\mu} = \tilde{\boldsymbol{\mu}}. \tag{56}$$

  for some $\tilde{\boldsymbol{\mu}}\in\mathbb{R}^d$
- Assumptions on the distribution of the randomization: $\mathbb{G}$ is continuous and has full support.

[22] Note that here we assume the affine selection region, written before as $\sqrt{n}\boldsymbol{A}_M\boldsymbol{T}+\boldsymbol{\omega}\in\boldsymbol{H}_M$, is given as $\sqrt{n}\boldsymbol{A}_M\boldsymbol{T}+\boldsymbol{\omega}\geqslant\boldsymbol{b}_M$ for a sequence of vectors $\boldsymbol{b}_M\in\mathbb{R}^d$. We do this in order to simplify the notation for the convergence of the selection region.

Theorem 13 (Selective CLT, version 1) *Assuming the conditions above plus* (CLT)*, the following holds*

$$\lim_{n\to\infty} \sup_{\mathbb{F}_n \in \mathcal{F}_n} \sup_{t\in[0,1]} \left| \mathbb{F}^* \left\{ \mathcal{P}\left(\sqrt{n}(\boldsymbol{T}-\boldsymbol{\mu}), \boldsymbol{A}^T\boldsymbol{\mu}, \sqrt{n}\boldsymbol{A}_M\boldsymbol{\mu}, \boldsymbol{A}_M, \boldsymbol{b}_M\right) \leqslant t \right\} - t \right| = 0$$

Proof

Denote $\Phi^*_{\boldsymbol{\Sigma}}$ the distribution of $\boldsymbol{G} \sim \mathcal{N}_p(\boldsymbol{0}, \boldsymbol{\Sigma}) := \Phi_{\boldsymbol{\Sigma}}$ conditional on the selection

$$\widetilde{\boldsymbol{A}}\boldsymbol{G} + \tilde{\boldsymbol{\mu}} + \boldsymbol{\omega} \geqslant \tilde{\boldsymbol{b}}.$$

We know that

$$\Phi^*_{\boldsymbol{\Sigma}}\left\{ \mathcal{P}(\boldsymbol{G}, \boldsymbol{A}^T\boldsymbol{\mu}, \tilde{\boldsymbol{\mu}}, \widetilde{\boldsymbol{A}}, \tilde{\boldsymbol{b}}) \leqslant t \right\} = t.$$

Suffices to show that for every continuous bounded function $g : \mathbb{R} \to \mathbb{R}$, we have

$$\mathbb{E}_{\mathbb{F}^*_n}\left[g\left(\mathcal{P}\left(\sqrt{n}(\boldsymbol{T}-\boldsymbol{\mu}), \boldsymbol{A}^T\boldsymbol{\mu}, \sqrt{n}\boldsymbol{A}_M\boldsymbol{\mu}, \boldsymbol{A}_M, \boldsymbol{b}_M\right)\right)\right] \to \mathbb{E}_{\Phi^*_{\boldsymbol{\Sigma}}}\left[g(\mathcal{P}\left(\boldsymbol{G}, \boldsymbol{A}^T\boldsymbol{\mu}, \tilde{\boldsymbol{\mu}}, \widetilde{\boldsymbol{A}}, \tilde{\boldsymbol{b}}\right)\right]$$

as $n \to \infty$. Denoting the selective likelihood ratios for the data originally distributed as $\boldsymbol{T} \sim \mathbb{F}_n$ as

$$\ell_{\mathbb{F}_n}(\sqrt{n}(\boldsymbol{T}-\boldsymbol{\mu}), \sqrt{n}\boldsymbol{A}_M\boldsymbol{\mu}, \boldsymbol{A}_M, \boldsymbol{b}_M) = \frac{\mathbb{G}\left\{\sqrt{n}\boldsymbol{A}_M(\boldsymbol{T}-\boldsymbol{\mu}) + \sqrt{n}\boldsymbol{A}_M\boldsymbol{\mu} + \boldsymbol{\omega} \geqslant \boldsymbol{b}_M\right\}}{\mathbb{E}_{\mathbb{F}_n}\left[\mathbb{G}\left\{\boldsymbol{\omega} \geqslant \boldsymbol{b}_M - \sqrt{n}\boldsymbol{A}_M(\boldsymbol{T}'-\boldsymbol{\mu}) - \sqrt{n}\boldsymbol{A}_M\boldsymbol{\mu}\right\}\right]}$$

and Gaussian one as

$$\ell_{\Phi_{\boldsymbol{\Sigma}}}(\boldsymbol{G}, \tilde{\boldsymbol{\mu}}, \widetilde{\boldsymbol{A}}, \tilde{\boldsymbol{b}}) = \frac{\mathbb{G}\left\{\widetilde{\boldsymbol{A}}\boldsymbol{G} + \tilde{\boldsymbol{\mu}} + \boldsymbol{\omega} \geqslant \tilde{\boldsymbol{b}}\right\}}{\mathbb{E}_{\Phi_{\boldsymbol{\Sigma}}}\left[\mathbb{G}\left\{\boldsymbol{\omega} \geqslant \tilde{\boldsymbol{b}} - \widetilde{\boldsymbol{A}}\boldsymbol{G}' - \tilde{\boldsymbol{\mu}}\right\}\right]}$$

it suffices to show the following holds

$$\begin{aligned}&\mathbb{E}_{\mathbb{F}_n}\left[g\left(\mathcal{P}\left(\sqrt{n}(\boldsymbol{T}-\boldsymbol{\mu}), \boldsymbol{A}^T\boldsymbol{\mu}, \sqrt{n}\boldsymbol{A}_M\boldsymbol{\mu}, \boldsymbol{A}_M, \boldsymbol{b}_M\right)\right) \ell_{\mathbb{F}_n}(\sqrt{n}(\boldsymbol{T}-\boldsymbol{\mu}), \sqrt{n}\boldsymbol{A}_M\boldsymbol{\mu}, \boldsymbol{A}_M, \boldsymbol{b}_M)\right] \\ &\to \mathbb{E}_{\Phi_{\boldsymbol{\Sigma}}}\left[g\left(\mathcal{P}(\boldsymbol{G}, \boldsymbol{A}^T\boldsymbol{\mu}, \tilde{\boldsymbol{\mu}}, \widetilde{\boldsymbol{A}}, \tilde{\boldsymbol{b}})\right) \ell_{\Phi_{\boldsymbol{\Sigma}}}(\boldsymbol{G}, \tilde{\boldsymbol{\mu}}, \widetilde{\boldsymbol{A}}, \tilde{\boldsymbol{b}})\right]\end{aligned}$$

as $n \to \infty$ for an arbitrary sequence $\mathbb{F}_n \in \mathcal{F}_n$.

(CLT) assumption and continuity of $\mathbb{G}$ implies for all $\boldsymbol{t} \in \mathbb{R}^p$

$$\lim_{n\to\infty} \sup_{\boldsymbol{t}\in\mathbb{R}^p} \left| \ell_{\mathbb{F}_n}(\boldsymbol{t}, \sqrt{n}\boldsymbol{A}_M\boldsymbol{\mu}, \boldsymbol{A}_M, \boldsymbol{b}_M) - \ell_{\Phi_{\boldsymbol{\Sigma}}}(\boldsymbol{t}, \widetilde{\boldsymbol{A}}, \tilde{\boldsymbol{\mu}}, \widetilde{\boldsymbol{A}}, \tilde{\boldsymbol{b}}) \right| = 0.$$

This implies (by Lebesgue dominated convergence theorem)

$$\lim_{n\to\infty} \mathbb{E}_{\mathbb{F}_n}\left[\left| \ell_{\mathbb{F}_n}\left(\sqrt{n}(\boldsymbol{T}-\boldsymbol{\mu}), \sqrt{n}\boldsymbol{A}_M\boldsymbol{\mu}, \boldsymbol{A}_M, \boldsymbol{b}_M\right) - \ell_{\Phi_{\boldsymbol{\Sigma}}}(\sqrt{n}(\boldsymbol{T}-\boldsymbol{\mu}), \tilde{\boldsymbol{\mu}}, \widetilde{\boldsymbol{A}}, \tilde{\boldsymbol{b}}) \right| \right] = 0.$$

Since g is bounded function, it suffices to show

$$\begin{aligned}&\mathbb{E}_{\mathbb{F}_n}\left[g\left(\mathcal{P}\left(\sqrt{n}(\boldsymbol{T}-\boldsymbol{\mu}), \boldsymbol{A}^T\boldsymbol{\mu}, \sqrt{n}\boldsymbol{A}_M\boldsymbol{\mu}, \boldsymbol{A}_M, \boldsymbol{b}_M\right)\right) \ell_{\Phi_{\boldsymbol{\Sigma}}}(\sqrt{n}(\boldsymbol{T}-\boldsymbol{\mu}), \tilde{\boldsymbol{\mu}}, \widetilde{\boldsymbol{A}}, \tilde{\boldsymbol{b}})\right] \\ &\to \mathbb{E}_{\Phi_{\boldsymbol{\Sigma}}}\left[g\left(\mathcal{P}(\boldsymbol{G}, \boldsymbol{A}^T\boldsymbol{\mu}, \tilde{\boldsymbol{\mu}}, \widetilde{\boldsymbol{A}}, \tilde{\boldsymbol{b}})\right) \ell_{\Phi_{\boldsymbol{\Sigma}}}(\boldsymbol{G}, \tilde{\boldsymbol{\mu}}, \widetilde{\boldsymbol{A}}, \tilde{\boldsymbol{b}})\right]\end{aligned}$$

as $n \to \infty$.

By the uniform Continuous mapping theorem the convergence above holds in distribution. To go to expectation, it suffices to have g and $\ell_{\Phi_{\boldsymbol{\Sigma}}}$ to be bounded. g is bounded by assumption and $\ell_{\Phi_{\boldsymbol{\Sigma}}}$ is bounded above as long as $\mathbb{E}_{\Phi_{\boldsymbol{\Sigma}}}\left[\mathbb{G}\left\{\boldsymbol{\omega} : \boldsymbol{\omega} \geqslant \tilde{\boldsymbol{b}} - \widetilde{\boldsymbol{A}}\boldsymbol{G} - \tilde{\boldsymbol{\mu}}\right\}\right] > 0$, which holds since $\mathbb{G}$ has full support by assumption.

■

B. Proving selective CLT for Lipschitz randomization

B.1. Smoothness of the Gaussian likelihood and the pivot for Lipschitz randomization

Recall that

$$Q_1(\boldsymbol{z};\boldsymbol{\Delta}) = \mathbb{G}\{\boldsymbol{A}_M\boldsymbol{z} + \boldsymbol{A}_M\boldsymbol{\Delta} + \boldsymbol{\omega} \in \boldsymbol{H}_M\},$$

and the Gaussian selective likelihood ratio becomes

$$\ell_{\Phi_n}(\boldsymbol{z}) = \frac{\mathbb{G}\{\boldsymbol{A}_M\boldsymbol{z} + \boldsymbol{A}_M\boldsymbol{\Delta} + \boldsymbol{\omega} \in \boldsymbol{H}_M\}}{\Phi_n \times \mathbb{G}\{\boldsymbol{A}_M\boldsymbol{Z} + \boldsymbol{A}_M\boldsymbol{\Delta} + \boldsymbol{\omega} \in \boldsymbol{H}_M\}} = \frac{Q_1(\boldsymbol{z};\boldsymbol{\Delta})}{\mathbb{E}_{\Phi_n}\left[Q_1(\boldsymbol{Z};\boldsymbol{\Delta})\right]}.$$

Lemma 14 (Lipschitz randomization: smoothness of the Gaussian likelihood) *Assuming* (NA) *and* (S) *hold, we have*

$$\partial^{\boldsymbol{\alpha}}_{\boldsymbol{z}}\ell_{\Phi_n}(\boldsymbol{z}) \leqslant \frac{C_M \cdot K_g}{\mathbb{E}_{\Phi_n}\left[e^{-K_g\|\boldsymbol{A}_M\boldsymbol{Z}\|_h}\right]} e^{K_g\|\boldsymbol{A}_M\boldsymbol{z}\|_h} = O(e^{K_g\|\boldsymbol{A}_M\boldsymbol{z}\|_h}).$$

Proof

Using the change of variables $\boldsymbol{\omega}' = \boldsymbol{\omega} + \boldsymbol{A}_M\boldsymbol{z} + \boldsymbol{A}_M\boldsymbol{\Delta}$ and the Lipschitz assumption on $\tilde{g}$, we have the lower bound

$$\begin{aligned}
Q_1(\boldsymbol{z};\boldsymbol{\Delta}) &= \int\limits_{\boldsymbol{\omega}\in\boldsymbol{H}_M - \boldsymbol{A}_M\boldsymbol{z} - \boldsymbol{A}_M\boldsymbol{\Delta}} \frac{1}{C_g}\exp(-\tilde{g}(\boldsymbol{\omega}))d\boldsymbol{\omega} \\
&= \frac{1}{C_g}\int\limits_{\boldsymbol{\omega}'\in\boldsymbol{H}_M} \exp\left(-\tilde{g}(\boldsymbol{\omega}' - \boldsymbol{A}_M\boldsymbol{z} - \boldsymbol{A}_M\boldsymbol{\Delta})\right) d\boldsymbol{\omega}' \\
&\geqslant \frac{1}{C_g}\int\limits_{\boldsymbol{\omega}'\in\boldsymbol{H}_M} \exp\left(-\tilde{g}(\boldsymbol{\omega}' - \boldsymbol{A}_M\boldsymbol{\Delta}) - K_g\|\boldsymbol{A}_M\boldsymbol{z}\|_h\right) d\boldsymbol{\omega}' \\
&= e^{-K_g\|\boldsymbol{A}_M\boldsymbol{z}\|_h}\int\limits_{\boldsymbol{\omega}'\in\boldsymbol{H}_M} \frac{1}{C_g}\exp\left(-\tilde{g}(\boldsymbol{\omega}' - \boldsymbol{A}_M\boldsymbol{\Delta})\right) d\boldsymbol{\omega}' \\
&= Q_1(\boldsymbol{0};\boldsymbol{\Delta})\cdot e^{-K_g\|\boldsymbol{A}_M\boldsymbol{z}\|_h},
\end{aligned} \tag{57}$$

implying

$$\mathbb{E}_{\Phi_n}\left[Q_1(\boldsymbol{Z};\boldsymbol{\Delta})\right] \geqslant Q_1(\boldsymbol{0};\boldsymbol{\Delta})\cdot\mathbb{E}_{\Phi_n}\left[e^{-K_g\|\boldsymbol{A}_M\boldsymbol{Z}\|_h}\right]. \tag{58}$$

Similarly, using the Lipschitz property on $\tilde{g}$ again we have an upper bound

$$\begin{aligned}
Q_1(\boldsymbol{z};\boldsymbol{\Delta}) &\leqslant \frac{1}{C_g}\int\limits_{\boldsymbol{\omega}'\in\boldsymbol{H}_M} \exp\left(-\tilde{g}(\boldsymbol{\omega}' - \boldsymbol{A}_M\boldsymbol{\Delta}) + K_g\|\boldsymbol{A}_M\boldsymbol{z}\|_h\right) d\boldsymbol{\omega}' \\
&= e^{K_g\|\boldsymbol{A}_M\boldsymbol{z}\|_h}\int\limits_{\boldsymbol{\omega}'\in\boldsymbol{H}_M} \frac{1}{C_g}\exp\left(-\tilde{g}(\boldsymbol{\omega}' - \boldsymbol{A}_M\boldsymbol{\Delta})\right) d\boldsymbol{\omega}' \\
&= Q_1(\boldsymbol{0};\boldsymbol{\Delta})\cdot e^{K_g\|\boldsymbol{A}_M\boldsymbol{z}\|_h}.
\end{aligned} \tag{59}$$

By the smoothness assumption on $\tilde{g}$, we have

$$\begin{aligned}
|\partial_{\boldsymbol{z}}^{\boldsymbol{\alpha}} Q_1(\boldsymbol{z};\boldsymbol{\Delta})| &= \frac{1}{C_g}\left|\int\limits_{\boldsymbol{\omega}'\in\boldsymbol{H}_M} \partial_{\boldsymbol{z}}^{\boldsymbol{\alpha}} \exp(-\tilde{g}(\boldsymbol{\omega}'-\boldsymbol{A}_M\boldsymbol{z}-\boldsymbol{A}_M\boldsymbol{\Delta}))d\boldsymbol{\omega}'\right| \\
&\leqslant C_M\cdot K_g \int\limits_{\boldsymbol{\omega}'\in\boldsymbol{H}_M} \frac{1}{C_g}\exp(-\tilde{g}(\boldsymbol{\omega}'-\boldsymbol{A}_M\boldsymbol{z}-\boldsymbol{A}_M\boldsymbol{\Delta}))d\boldsymbol{\omega}' \\
&= C_M\cdot K_g\cdot Q_1(\boldsymbol{z};\boldsymbol{\Delta}) \\
&\leqslant C_M\cdot K_g\cdot Q_1(\boldsymbol{0};\boldsymbol{\Delta})\cdot e^{K_g\|\boldsymbol{A}_M\boldsymbol{z}\|_h}, \qquad (60)
\end{aligned}$$

where C_M comes from differentiating $\boldsymbol{A}_M\boldsymbol{z}$ with respect to $\boldsymbol{z}$ and the last inequality follows by using the upper bound on $Q_1(\boldsymbol{z};\boldsymbol{\Delta})$ from (59). Combining (57) and (60) and we have

$$\partial_{\boldsymbol{z}}^{\boldsymbol{\alpha}}\ell_{\Phi_n}(\boldsymbol{z}) = \frac{\partial_{\boldsymbol{z}}^{\boldsymbol{\alpha}} Q_1(\boldsymbol{z};\boldsymbol{\Delta})}{\mathbb{E}_{\Phi_n} Q_1(\boldsymbol{Z};\boldsymbol{\Delta})} \leqslant \frac{C_M\cdot K_g}{\mathbb{E}_{\Phi_n}\left[e^{-K_g\|\boldsymbol{A}_M\boldsymbol{Z}\|_h}\right]} e^{K_g\|\boldsymbol{A}_M\boldsymbol{z}\|_h}.$$

■

Denote

$$Q_2(\boldsymbol{\tau};\boldsymbol{z}_{\boldsymbol{A}},\boldsymbol{\Delta}) = \mathbb{G}\{\boldsymbol{A}_M\boldsymbol{z}_{\boldsymbol{A}} + \boldsymbol{A}_M\boldsymbol{C}\boldsymbol{\tau} + \boldsymbol{A}_M\boldsymbol{\Delta} + \boldsymbol{\omega}\in\boldsymbol{H}_M\}$$

and

$$\begin{aligned}
Q_3(\boldsymbol{z}_{\boldsymbol{A}};\boldsymbol{\Delta}) &= \int\limits_{\mathbb{R}^a} \mathbb{G}\{\boldsymbol{A}_M\boldsymbol{z}_{\boldsymbol{A}} + \boldsymbol{A}_M\boldsymbol{C}\boldsymbol{\tau} + \boldsymbol{A}_M\boldsymbol{\Delta} + \boldsymbol{\omega}\in\boldsymbol{H}_M\}\phi_a(\boldsymbol{\tau};\boldsymbol{0},\boldsymbol{\Sigma}_{\boldsymbol{A}})d\boldsymbol{\tau} \\
&= \mathbb{E}_{\boldsymbol{G}\sim\mathcal{N}_a(\boldsymbol{0},\boldsymbol{\Sigma}_{\boldsymbol{A}})}\left[Q_2(\boldsymbol{G};\boldsymbol{z}_{\boldsymbol{A}},\boldsymbol{\Delta})\right].
\end{aligned}$$

Then the pivot equals

$$\mathcal{P}(\boldsymbol{Z},\boldsymbol{A}_M\boldsymbol{\Delta},\boldsymbol{A}_M,\boldsymbol{H}_M) = \frac{\int\limits_{\|\boldsymbol{\Sigma}_{\boldsymbol{A}}^{-1/2}\boldsymbol{\tau}\|_2\geqslant\|\boldsymbol{\Sigma}_{\boldsymbol{A}}^{-1/2}\boldsymbol{A}^T\boldsymbol{Z}\|_2} Q_2(\boldsymbol{\tau};\boldsymbol{Z}_{\boldsymbol{A}},\boldsymbol{\Delta})\phi_a(\boldsymbol{\tau};\boldsymbol{0},\boldsymbol{\Sigma}_{\boldsymbol{A}})d\boldsymbol{\tau}}{Q_3(\boldsymbol{Z}_{\boldsymbol{A}};\boldsymbol{\Delta})},$$

where with the abuse of notation we write different the arguments of the pivot slightly differently in this section.

Lemma 15 (Lipschitz randomization: smoothness of the pivot) *Assuming* (NA) *and* (S) *hold, we have*

$$\begin{aligned}
|\partial_{\boldsymbol{z}}^{\boldsymbol{\alpha}}\mathcal{P}(\boldsymbol{z},\boldsymbol{A}_M\boldsymbol{\Delta},\boldsymbol{A}_M,\boldsymbol{H}_M)| &\leqslant C(\boldsymbol{A},\boldsymbol{A}_M,K_g) + C(\boldsymbol{A})\frac{e^{K_g\|\boldsymbol{A}_M\boldsymbol{C}\boldsymbol{\Sigma}_{\boldsymbol{A}}\|_{h,2}\|\boldsymbol{\Sigma}_{\boldsymbol{A}}^{-1/2}\boldsymbol{A}^T\boldsymbol{z}\|_2}}{\mathbb{E}_{\boldsymbol{G}\sim\mathcal{N}_a(\boldsymbol{0},\boldsymbol{\Sigma}_{\boldsymbol{A}})}\left[e^{-K_g\|\boldsymbol{A}_M\boldsymbol{C}\boldsymbol{G}\|_h}\right]} \\
&= O\left(1+e^{K_g\|\boldsymbol{A}_M\boldsymbol{C}\boldsymbol{\Sigma}_{\boldsymbol{A}}\|_{h,2}\|\boldsymbol{\Sigma}_{\boldsymbol{A}}^{-1/2}\boldsymbol{A}^T\boldsymbol{z}\|_2}\right),
\end{aligned}$$

where the constants above depend only on their arguments.

Proof

By the change of variables $\boldsymbol{\omega} = \boldsymbol{\omega}' - \boldsymbol{A}_M \boldsymbol{z}_{\boldsymbol{A}} - \boldsymbol{A}_M \boldsymbol{C}\boldsymbol{\tau} - \boldsymbol{A}_M \boldsymbol{\Delta}$ and using the Lipschitz property of the randomization we get the lower bound on $Q_2(\boldsymbol{\tau}; \boldsymbol{z}_{\boldsymbol{A}}, \boldsymbol{\Delta})$:

$$\begin{aligned} Q_2(\boldsymbol{\tau}; \boldsymbol{z}_{\boldsymbol{A}}, \boldsymbol{\Delta}) &= \int\limits_{\boldsymbol{\omega} \in \boldsymbol{H}_M - \boldsymbol{A}_M \boldsymbol{z}_{\boldsymbol{A}} - \boldsymbol{A}_M \boldsymbol{C}\boldsymbol{\tau} - \boldsymbol{A}_M \boldsymbol{\Delta}} \frac{1}{C_g} \exp(-\tilde{g}(\boldsymbol{\omega})) d\boldsymbol{\omega} \\ &= \frac{1}{C_g} \int\limits_{\boldsymbol{\omega}' \in \boldsymbol{H}_M} \exp\left(-\tilde{g}(\boldsymbol{\omega}' - \boldsymbol{A}_M \boldsymbol{z}_{\boldsymbol{A}} - \boldsymbol{A}_M \boldsymbol{C}\boldsymbol{\tau} - \boldsymbol{A}_M \boldsymbol{\Delta})\right) d\boldsymbol{\omega}' \\ &\geqslant \frac{1}{C_g} \int\limits_{\boldsymbol{\omega}' \in \boldsymbol{H}_M} \exp\left(-\tilde{g}(\boldsymbol{\omega}' - \boldsymbol{A}_M \boldsymbol{z}_{\boldsymbol{A}} - \boldsymbol{A}_M \boldsymbol{\Delta}) - K_g \|\boldsymbol{A}_M \boldsymbol{C}\boldsymbol{\tau}\|_h\right) d\boldsymbol{\omega}' \\ &= Q_2(\boldsymbol{0}; \boldsymbol{z}_{\boldsymbol{A}}, \boldsymbol{\Delta}) \cdot e^{-K_g \|\boldsymbol{A}_M \boldsymbol{C}\boldsymbol{\tau}\|_h}. \end{aligned}$$

Hence

$$Q_3(\boldsymbol{z}_{\boldsymbol{A}}; \boldsymbol{\Delta}) \geqslant Q_2(\boldsymbol{0}; \boldsymbol{z}_{\boldsymbol{A}}, \boldsymbol{\Delta}) \cdot \mathbb{E}_{\boldsymbol{G} \sim \mathcal{N}_a(\boldsymbol{0}, \boldsymbol{\Sigma}_{\boldsymbol{A}})} \left[e^{-K_g \|\boldsymbol{A}_M \boldsymbol{C}\boldsymbol{G}\|_h} \right]. \tag{61}$$

As for the upper bound on $Q_2(\boldsymbol{\tau}; \boldsymbol{z}_{\boldsymbol{A}}, \boldsymbol{\Delta})$, we have

$$\begin{aligned} Q_2(\boldsymbol{\tau}; \boldsymbol{z}_{\boldsymbol{A}}, \boldsymbol{\Delta}) &\leqslant \frac{1}{C_g} \int\limits_{\boldsymbol{\omega}' \in \boldsymbol{H}_M} \exp\left(-\tilde{g}(\boldsymbol{\omega}' - \boldsymbol{A}_M \boldsymbol{z}_{\boldsymbol{A}} - \boldsymbol{A}_M \boldsymbol{\Delta}) + K_g \|\boldsymbol{A}_M \boldsymbol{C}\boldsymbol{\tau}\|\right) d\boldsymbol{\omega}' \\ &= Q_2(\boldsymbol{0}; \boldsymbol{z}_{\boldsymbol{A}}; \boldsymbol{\Delta}) \cdot e^{K_g \|\boldsymbol{A}_M \boldsymbol{C}\boldsymbol{\tau}\|_h}. \end{aligned}$$

Decomposing $\boldsymbol{\tau} = \boldsymbol{\Sigma}_{\boldsymbol{A}}^{1/2} \|\boldsymbol{\Sigma}_A^{-1/2} \boldsymbol{\tau}\|_2 \frac{\boldsymbol{\Sigma}_{\boldsymbol{A}}^{-1/2} \boldsymbol{\tau}}{\|\boldsymbol{\Sigma}_A^{-1/2} \boldsymbol{\tau}\|_2}$ and denoting $r = \|\boldsymbol{\Sigma}_A^{-1/2} \boldsymbol{\tau}\|_2$ and $\boldsymbol{u} = \frac{\boldsymbol{\Sigma}_{\boldsymbol{A}}^{-1/2} \boldsymbol{\tau}}{\|\boldsymbol{\Sigma}_A^{-1/2} \boldsymbol{\tau}\|_2}$, we know that under $\boldsymbol{\tau} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{\Sigma}_{\boldsymbol{A}})$, r and $\boldsymbol{u}$ are independent random variables. Denote their densities with $f_r(r)$ and $f_{\boldsymbol{u}} e(\boldsymbol{u})$, respectively.

Then the derivative of the numerator with respect to $\boldsymbol{A}^T \boldsymbol{z}$ satisfies

$$\begin{aligned} &\left| \partial^{\boldsymbol{\alpha}}_{\boldsymbol{A}^T \boldsymbol{z}} \int_{r \geqslant \|\boldsymbol{\Sigma}_{\boldsymbol{A}}^{-1/2} \boldsymbol{A}^T \boldsymbol{z}\|_2} \int_{\boldsymbol{u} \in \mathcal{S}^a} Q_2\left(r \boldsymbol{\Sigma}_{\boldsymbol{A}}^{1/2} \boldsymbol{u}; \boldsymbol{z}_{\boldsymbol{A}}, \boldsymbol{\Delta}\right) f_r(r) f_{\boldsymbol{u}}(\boldsymbol{u}) dr d\boldsymbol{u} \right| \\ &\leqslant C(\boldsymbol{A}) \cdot \int_{\boldsymbol{u} \in \mathcal{S}^a} Q_2\left(\|\boldsymbol{\Sigma}_{\boldsymbol{A}}^{-1/2} \boldsymbol{A}^T \boldsymbol{z}\|_2 \boldsymbol{\Sigma}_{\boldsymbol{A}}^{1/2} \boldsymbol{u}; \boldsymbol{z}_{\boldsymbol{A}}, \boldsymbol{\Delta}\right) f_r(\|\boldsymbol{\Sigma}_{\boldsymbol{A}}^{-1/2} \boldsymbol{A}^T \boldsymbol{z}\|_2) f_{\boldsymbol{u}}(\boldsymbol{u}) d\boldsymbol{u} \\ &\leqslant C(\boldsymbol{A}) \cdot Q_2(\boldsymbol{0}; \boldsymbol{z}_{\boldsymbol{A}}, \boldsymbol{\Delta}) \cdot e^{K_g \|\boldsymbol{A}_M \boldsymbol{C} \boldsymbol{\Sigma}_{\boldsymbol{A}}^{1/2} \boldsymbol{u}\|_h \|\boldsymbol{\Sigma}_{\boldsymbol{A}}^{-1/2} \boldsymbol{A}^T \boldsymbol{z}\|_2} f_r(\|\boldsymbol{\Sigma}_{\boldsymbol{A}}^{-1/2} \boldsymbol{A}^T \boldsymbol{z}\|_2) \int_{\boldsymbol{u} \in \mathcal{S}^a} f_{\boldsymbol{u}}(\boldsymbol{u}) d\boldsymbol{u}, \end{aligned}$$

where $\mathcal{S}^a$ is a unit sphere in $\mathbb{R}^a$ and in the last inequality we used the upper bound on Q_2. Since the last integral is 1 and the density $f_r(\cdot)$ is uniformly bounded by a constant, we get that the derivative above is upper bounded by

$$C(\boldsymbol{A}) \cdot Q_2(\boldsymbol{0}; \boldsymbol{z}_{\boldsymbol{A}}; \boldsymbol{\Delta}) \cdot e^{K_g \|\boldsymbol{A}_M \boldsymbol{C} \boldsymbol{\Sigma}_{\boldsymbol{A}}\|_{h,2} \|\boldsymbol{\Sigma}_{\boldsymbol{A}}^{-1/2} \boldsymbol{A}^T \boldsymbol{z}\|_2}.$$

Combing this upper bound with the lower bound on the denominator from (61), we get that the derivative of the pivot with respect to $\boldsymbol{A}^T \boldsymbol{z}$ is upper bounded by

$$|\partial^{\boldsymbol{\alpha}}_{\boldsymbol{A}^T \boldsymbol{z}} \mathcal{P}(\boldsymbol{z}, \boldsymbol{A}_M \boldsymbol{\Delta}, \boldsymbol{A}_M, \boldsymbol{H}_M)| \leqslant C(\boldsymbol{A}) \frac{e^{K_g \|\boldsymbol{A}_M \boldsymbol{C} \boldsymbol{\Sigma}_{\boldsymbol{A}}\|_{h,2} \|\boldsymbol{\Sigma}_{\boldsymbol{A}}^{-1/2} \boldsymbol{A}^T \boldsymbol{z}\|_2}}{\mathbb{E}_{\boldsymbol{G} \sim \mathcal{N}_a(\boldsymbol{0}, \boldsymbol{\Sigma}_{\boldsymbol{A}})} \left[e^{-K_g \|\boldsymbol{A}_M \boldsymbol{C}\boldsymbol{G}\|_h} \right]}.$$

Now we turn to the derivative of the pivot with respect to $\boldsymbol{z_A}$. The derivative of the integrand of the numerator of the pivot is upper bounded by

$$\begin{aligned}
|\partial_{\boldsymbol{z}}^{\boldsymbol{\alpha}} Q_2(\boldsymbol{\tau}; \boldsymbol{z_A}, \boldsymbol{\Delta})| &= \frac{1}{C_g} \left| \int\limits_{\boldsymbol{\omega}' \in \boldsymbol{H}_M} \partial_{\boldsymbol{z}}^{\boldsymbol{\alpha}} \exp\left(-\tilde{g}(\boldsymbol{\omega}' - \boldsymbol{A}_M \boldsymbol{z_A} - \boldsymbol{A}_M \boldsymbol{C}\boldsymbol{\tau} - \boldsymbol{A}_M \boldsymbol{\Delta})\right) d\boldsymbol{\omega}' \right| \\
&\leqslant C(\boldsymbol{A}_M, \boldsymbol{A}) \frac{K_g}{C_g} \int\limits_{\boldsymbol{\omega}' \in \boldsymbol{H}_M} \exp\left(-\tilde{g}(\boldsymbol{\omega}' - \boldsymbol{A}_M \boldsymbol{z_A} - \boldsymbol{A}_M \boldsymbol{C}\boldsymbol{\tau} - \boldsymbol{A}_M \boldsymbol{\Delta})\right) d\boldsymbol{\omega}' \\
&= C(\boldsymbol{A}_M, \boldsymbol{A}) \cdot K_g \cdot Q_2(\boldsymbol{\tau}; \boldsymbol{z_A}, \boldsymbol{\Delta})
\end{aligned}$$

and, similarly, the derivative of the denominator is upper bounded by

$$\left|\partial_{\boldsymbol{z_A}}^{\boldsymbol{\alpha}} Q_3(\boldsymbol{z}_A; \boldsymbol{\Delta})\right| \leqslant C(\boldsymbol{A}_M, \boldsymbol{A}) \cdot K_g \cdot Q_3(\boldsymbol{z_A}; \boldsymbol{\Delta}). \tag{62}$$

Note that the derivative of the pivot with respect to $\boldsymbol{z_A}$ will be the sum of the term of the form

$$\frac{\int\limits_{\|\boldsymbol{\Sigma}_{\boldsymbol{A}}^{-1/2}\boldsymbol{\tau}\|_2 \geqslant \|\boldsymbol{\Sigma}_{\boldsymbol{A}}^{-1/2}\boldsymbol{A}^T\boldsymbol{z}\|_2} \partial_{\boldsymbol{z}}^{\boldsymbol{\alpha}} Q_2(\boldsymbol{\tau}; \boldsymbol{z_A}, \boldsymbol{\Delta}) \phi_a(\boldsymbol{\tau}; \boldsymbol{0}, \boldsymbol{\Sigma_A}) d\boldsymbol{\tau}}{Q_3(\boldsymbol{z_A}; \boldsymbol{\Delta})}$$

and

$$\frac{\left(\int\limits_{\|\boldsymbol{\Sigma}_{\boldsymbol{A}}^{-1/2}\boldsymbol{\tau}\|_2 \geqslant \|\boldsymbol{\Sigma}_{\boldsymbol{A}}^{-1/2}\boldsymbol{A}^T\boldsymbol{z}\|_2} Q_2(\boldsymbol{\tau}; \boldsymbol{z_A}, \boldsymbol{\Delta}) \phi_a(\boldsymbol{\tau}; \boldsymbol{0}, \boldsymbol{\Sigma_A}) d\boldsymbol{\tau}\right) \cdot \partial_{\boldsymbol{z_A}}^{\boldsymbol{\alpha}} Q_3(\boldsymbol{z_A}; \boldsymbol{\Delta})}{Q_3(\boldsymbol{z_A}; \boldsymbol{\Delta})^2}$$

The first term above in the absolute value is upper bounded by

$$\begin{aligned}
\frac{\int\limits_{\mathbb{R}^a} |\partial_{\boldsymbol{z}}^{\boldsymbol{\alpha}} Q_2(\boldsymbol{\tau}; \boldsymbol{z_A}, \boldsymbol{\Delta})| \phi_a(\boldsymbol{\tau}; \boldsymbol{0}, \boldsymbol{\Sigma_A}) d\boldsymbol{\tau}}{Q_3(\boldsymbol{z_A}; \boldsymbol{\Delta})} &\leqslant \frac{C(\boldsymbol{A}_M, \boldsymbol{A}) \cdot K_g \cdot \mathbb{E}_{\boldsymbol{G} \sim \mathcal{N}_a(\boldsymbol{0}, \boldsymbol{\Sigma_A})} \left[Q_2(\boldsymbol{\tau}; \boldsymbol{z_A}, \boldsymbol{\Delta})\right]}{Q_3(\boldsymbol{z_A}; \boldsymbol{\Delta})} \\
&= C(\boldsymbol{A}_M, \boldsymbol{A}) \cdot K_g
\end{aligned}$$

and the second term in absolute value is upper bounded by

$$\frac{\partial_{\boldsymbol{z_A}}^{\boldsymbol{\alpha}} Q_3(\boldsymbol{z_A}; \boldsymbol{\Delta})}{Q_3(\boldsymbol{z_A}; \boldsymbol{\Delta})} \leqslant C(\boldsymbol{A}_M, \boldsymbol{A}) \cdot K_g.$$

■

B.2. Applying Chatterjee's theorem for Lipschitz randomization

Denote as $\boldsymbol{y}_i$ the centered versions of $\boldsymbol{\xi}_i$, $i = 1, \ldots, n$. Let $\mathcal{M} : \mathbb{R}^{n \times p} \to \mathbb{R}^p$ denote the normalizing operator:

$$\mathcal{M}(\boldsymbol{y}_1, \ldots, \boldsymbol{y}_n) = \frac{1}{\sqrt{n}} \sum_{i=1}^{n} \boldsymbol{y}_i.$$

Then our normalized test statistic becomes $\boldsymbol{Z} = \mathcal{M}(\boldsymbol{y}_1, \ldots, \boldsymbol{y}_n)$. Also, take $\boldsymbol{g}_1, \ldots, \boldsymbol{g}_n \overset{i.i.d.}{\sim} \mathcal{N}_p(\boldsymbol{0}, \boldsymbol{\Sigma})$:

$$\boldsymbol{G} = \begin{pmatrix} \boldsymbol{g}_1^T \\ \vdots \\ \boldsymbol{g}_n^T \end{pmatrix}_{n \times p}.$$

Denote two sequences of matrices $\{\boldsymbol{W}_i : i = 0, \ldots, n\}$, $\{\widetilde{\boldsymbol{W}}_i : i = 0, \ldots, n\}$ as

$$\boldsymbol{W}_i = \begin{pmatrix} \boldsymbol{g}_1^T \\ \vdots \\ \boldsymbol{g}_{i-1}^T \\ \boldsymbol{0} \\ \boldsymbol{y}_{i+1}^T \\ \vdots \\ \boldsymbol{y}_n^T \end{pmatrix}_{n\times p}, \quad \widetilde{\boldsymbol{W}}_i = \begin{pmatrix} \boldsymbol{g}_1^T \\ \vdots \\ \boldsymbol{g}_{i-1}^T \\ \boldsymbol{g}_i^T \\ \boldsymbol{y}_{i+1}^T \\ \vdots \\ \boldsymbol{y}_n^T \end{pmatrix}_{n\times p}.$$

Here we present a generalization of Theorem 1 in [Chatterjee, 2005], which provides a bound on the difference between smooth functions of Gaussian and averages of i.i.d. random variables coming from a nonparametric distribution.

Let $\Omega_n : \mathbb{R}^p \to \mathbb{R}$ (we will later take $\Omega_n(\boldsymbol{z}) = \|\boldsymbol{A}_M \boldsymbol{z}\|_h$) be a sequence of norm operators and let $\mathcal{W}_n : \mathbb{R} \to \mathbb{R}$ be a sequence of weights (we will later take $\mathcal{W}_n(x) = e^{K_g x}$). For $f : \mathbb{R}^{n\times p} \to \mathbb{R}$, we define

$$\lambda_3^{(n)}(f) = \sup\left\{ \frac{\|\partial_{y_i}^3 f(\boldsymbol{y}_1, \ldots, \boldsymbol{y}_n)\|_\infty}{\mathcal{W}_n\left(\Omega_n(\mathcal{M}(\boldsymbol{y}_1, \ldots, \boldsymbol{y}_n))\right)}, \;\; i = 1, \ldots, n \right\}.$$

Now we have the notation necessary to state the theorem that bounds the difference between the expectations of a function applied to nonparametric and the Gaussian data in terms of the smoothness of the function.

Theorem 16 ([Chatterjee, 2005]) *For $\mathcal{W}_n$ increasing and convex, we have*

$$\begin{aligned} |\mathbb{E}_{\mathbb{F}_n}[f(\boldsymbol{Y})] - \mathbb{E}_{\Phi_n}[f(\boldsymbol{G})]| \leqslant & \frac{1}{2}\sum_{i=1}^n \lambda_3^{(n)}(f)\mathbb{E}\left[\mathcal{W}_n\left(2\cdot\Omega(\mathcal{M}(\boldsymbol{W}_i))\right)\right]\left(\mathbb{E}\left[\|\boldsymbol{g}_i\|_1^3\right] + \mathbb{E}\left[\|\boldsymbol{y}_i\|_1^3\right]\right) \\ & + \frac{1}{2}\sum_{i=1}^n \lambda_3^{(n)}(f)\mathbb{E}\left[\mathcal{W}_n\left(2\cdot\Omega(\mathcal{M}(\widetilde{\boldsymbol{W}}_i - \boldsymbol{W}_i))\right)\|\boldsymbol{g}_i\|_1^3\right] \\ & + \frac{1}{2}\sum_{i=1}^n \lambda_3^{(n)}(f)\mathbb{E}\left[\mathcal{W}_n\left(2\cdot\Omega(\mathcal{M}(\widetilde{\boldsymbol{W}}_{i-1} - \boldsymbol{W}_i))\right)\|\boldsymbol{y}_i\|_1^3\right]. \end{aligned} \tag{63}$$

We now apply Theorem 16 to show the Gaussian and non-parametric likelihoods are close in expectation and to show that the constructed pivot under the non-parametric distribution asymptotically behaves as if the data was Gaussian. Since we know that the pivot is uniform under the Gaussian distribution, it will be asymptotically uniform under a non-parametric distribution.

Lemma 17 (Lipschitz randomization: closeness of Gaussian and non-parametric LR) *Assuming (AL), (MGF), (NA) and (S) hold, we have*

$$\mathbb{E}_{\mathbb{F}_n}\left[|\ell_{\mathbb{F}_n}(\boldsymbol{Z}) - \ell_{\Phi_n}(\boldsymbol{Z})|\right] \leqslant \frac{1}{\sqrt{n}} C(C_\tau, C_M, K_g, K_G, \boldsymbol{\Sigma}),$$

where the constant on the RHS depends only on its arguments.

Proof

Since

$$\begin{aligned} \mathbb{E}_{\mathbb{F}_n}\left[|\ell_{\mathbb{F}_n}(\boldsymbol{Z}) - \ell_{\Phi_n}(\boldsymbol{Z})|\right] &= \mathbb{E}_{\mathbb{F}_n}\left[Q_1(\boldsymbol{Z};\boldsymbol{\Delta})\right]\left|\frac{1}{\mathbb{E}_{\mathbb{F}_n}\left[Q_1(\boldsymbol{Z};\boldsymbol{\Delta})\right]} - \frac{1}{\mathbb{E}_{\Phi_n}\left[Q_1(\boldsymbol{Z};\boldsymbol{\Delta})\right]}\right| \\ &= \frac{|\mathbb{E}_{\mathbb{F}_n}\left[Q_1(\boldsymbol{Z};\boldsymbol{\Delta})\right] - \mathbb{E}_{\Phi_n}\left[Q_1(\boldsymbol{Z};\boldsymbol{\Delta})\right]|}{\mathbb{E}_{\Phi_n}\left[Q_1(\boldsymbol{Z};\boldsymbol{\Delta})\right]}, \end{aligned}$$

we have

$$\mathbb{E}_{\mathbb{F}_n}\left[|\ell_{\mathbb{F}_n}(\boldsymbol{Z})-\ell_{\Phi_n}(\boldsymbol{Z})|\right] \leqslant \frac{\left|\mathbb{E}_{\mathbb{F}_n}\left[\widetilde{Q}_1(\boldsymbol{Z};\boldsymbol{\Delta})\right]-\mathbb{E}_{\Phi_n}\left[\widetilde{Q}_1(\boldsymbol{Z};\boldsymbol{\Delta})\right]\right|}{\mathbb{E}_{\Phi_n}\left[e^{-K_g\|\boldsymbol{A}_M\boldsymbol{Z}\|_h}\right]},$$

where we used the lower bound on $\mathbb{E}_{\Phi_n}\left[Q_1(\boldsymbol{Z};\boldsymbol{\Delta})\right]$ from (58) and $\widetilde{Q}_1(\boldsymbol{z};\boldsymbol{\Delta})$ denotes

$$\widetilde{Q}_1(\boldsymbol{z};\boldsymbol{\Delta}) = \frac{Q_1(\boldsymbol{z};\boldsymbol{\Delta})}{Q_1(\boldsymbol{0},\boldsymbol{\Delta})}.$$

Hence it suffices to provide an upper bound on

$$\left|\mathbb{E}_{\mathbb{F}_n}\left[\widetilde{Q}_1(\boldsymbol{Z};\boldsymbol{\Delta})\right]-\mathbb{E}_{\Phi_n}\left[\widetilde{Q}_1(\boldsymbol{Z};\boldsymbol{\Delta})\right]\right|.$$

We use Chatterjee technique here, providing an upper bound on the quantity above using the smoothness of $\widetilde{Q}_1(\boldsymbol{z};\boldsymbol{\Delta})$. From (60) we have

$$\left|\partial^{\boldsymbol{\alpha}}_{\boldsymbol{z}}\widetilde{Q}_1(\boldsymbol{z};\boldsymbol{\Delta})\right| \leqslant C_M\cdot K_g\cdot e^{K_g\|\boldsymbol{A}_M\boldsymbol{z}\|_h},$$

implying

$$\lambda_3^{(n)}(\widetilde{Q}_1) \leqslant \frac{1}{n^{3/2}}C_M\cdot K_g$$

for $\Omega_n(\boldsymbol{z}) = \|\boldsymbol{A}_M\boldsymbol{z}\|_h$, $\boldsymbol{z}\in\mathbb{R}^p$, and $\mathcal{W}_n(x) = e^{K_g x}$, $x\in\mathbb{R}$. Using Theorem 16, we have

$$\begin{aligned}
&\left|\mathbb{E}_{\mathbb{F}_n}\left[\widetilde{Q}_1(\boldsymbol{Z};\boldsymbol{\Delta})\right]-\mathbb{E}_{\Phi_n}\left[\widetilde{Q}_1(\boldsymbol{Z};\boldsymbol{\Delta})\right]\right| \\
&\leqslant \frac{1}{2}\sum_{i=1}^n \lambda_3^{(n)}(\widetilde{Q}_1)\mathbb{E}\left[e^{2K_g\|\boldsymbol{A}_M\mathcal{M}(\boldsymbol{W}_i)\|_h}\right]\left(\mathbb{E}\left[\|\boldsymbol{g}_i\|_1^3\right]+\mathbb{E}\left[\|\boldsymbol{y}_i\|_1^3\right]\right) \\
&+\frac{1}{2}\sum_{i=1}^n \lambda_3^{(n)}(\widetilde{Q}_1)\mathbb{E}\left[e^{2K_g\|\boldsymbol{A}_M\mathcal{M}(\widetilde{\boldsymbol{W}}_i-\boldsymbol{W}_i)\|_h}\|\boldsymbol{g}_i\|_1^3\right] \\
&+\frac{1}{2}\sum_{i=1}^n \lambda_3^{(n)}(\widetilde{Q}_1)\mathbb{E}\left[e^{2K_g\|\boldsymbol{A}_M\mathcal{M}(\widetilde{\boldsymbol{W}}_{i-1}-\boldsymbol{W}_i)\|_h}\|\boldsymbol{y}_i\|_1^3\right].
\end{aligned} \tag{64}$$

Following the proof of Theorem 9 in [Tian and Taylor, 2015b], which uses (MGF) assumption, the sum in the RHS in (64) is of order $O(n^{-1/2})$.

■

Using the same proof as in Lemma 17, we bound the difference between the expectations of a smooth function applied to the pivots constructed for non-parametric and Gaussian data. The result is also under the conditional distributions of the data ($\mathbb{F}_n^*$ and Φ_n^* for the non-parametric and Gaussian respectively).

Theorem 18 (Lipschitz randomization: selective CLT) *Assume (AL), (MGF), (NA) and (S) hold. Given a function $\mathcal{H}:\mathbb{R}\to\mathbb{R}$ with uniformly bounded derivatives up to the third order (call this bound $K_{\mathcal{H}}$), we have*

$$\left|\mathbb{E}_{\mathbb{F}_n^*}\mathcal{H}(\mathcal{P}_n(\boldsymbol{Z}))-\mathbb{E}_{\Phi_n^*}\mathcal{H}(\mathcal{P}_n(\boldsymbol{G}))\right| \leqslant \frac{1}{\sqrt{n}}C(C_\tau, C_M, K_g, K_{\mathcal{H}}).$$

Proof Since

$$\begin{aligned}
&\left|\mathbb{E}_{\mathbb{F}_n^*}\mathcal{H}(\mathcal{P}_n(\boldsymbol{Z})) - \mathbb{E}_{\Phi_n^*}\mathcal{H}(\mathcal{P}_n(\boldsymbol{G}))\right| \\
&= \left|\mathbb{E}_{\mathbb{F}_n}\left[\mathcal{H}(\mathcal{P}_n(\boldsymbol{Z})\ell_{\mathbb{F}_n}(\boldsymbol{Z})\right] - \mathbb{E}_{\Phi_n}\left[\mathcal{H}(\mathcal{P}_n(\boldsymbol{G})\ell_{\Phi_n}(\boldsymbol{G})\right]\right|, \\
&= \left|\mathbb{E}_{\mathbb{F}_n}\left[\mathcal{H}(\mathcal{P}_n(\boldsymbol{Z})\left(\ell_{\mathbb{F}_n}(\boldsymbol{Z}) - \ell_{\Phi_n}(\boldsymbol{Z})\right)\right]\right| + \left|\mathbb{E}_{\mathbb{F}_n}\left[\mathcal{H}(\mathcal{P}_n(\boldsymbol{Z})\ell_{\Phi_n}(\boldsymbol{Z})\right] - \mathbb{E}_{\mathbb{F}_n}\left[\mathcal{H}(\mathcal{P}_n(\boldsymbol{G})\ell_{\Phi_n}(\boldsymbol{G})\right]\right| \\
&\leqslant K_{\mathcal{H}} \cdot \mathbb{E}_{\mathbb{F}_n}\left[\left|\ell_{\mathbb{F}_n}(\boldsymbol{Z}) - \ell_{\Phi_n}(\boldsymbol{Z})\right|\right] + \left|\mathbb{E}_{\mathbb{F}_n}\left[\mathcal{H}(\mathcal{P}_n(\boldsymbol{Z})\ell_{\Phi_n}(\boldsymbol{Z})\right] - \mathbb{E}_{\mathbb{F}_n}\left[\mathcal{H}(\mathcal{P}_n(\boldsymbol{G})\ell_{\Phi_n}(\boldsymbol{G})\right]\right|,
\end{aligned}$$

to bound the first term above we use the result of Lemma 17 and to bound the second term we apply Theorem 16 to

$$f(\boldsymbol{Y}) = \mathcal{H}(\mathcal{P}_n(\mathcal{M}(\boldsymbol{Y}))) \cdot \ell_{\Phi_n}(\mathcal{M}(\boldsymbol{Y})).$$

Applying Theorem 16 is then done in the same way as in Lemma 17, where we additionally use Lemma 15.

■

C. Proving Selective CLT for Gaussian randomization

C.1. Smoothness of the Gaussian likelihood and the pivot for Gaussian randomization

Denoting the probability of selection conditional on the data as

$$Q_1(\boldsymbol{z}; \boldsymbol{\Delta}) = \mathbb{G}\{\boldsymbol{A}_M\boldsymbol{z} + \boldsymbol{A}_M\boldsymbol{\Delta} + \boldsymbol{\omega} \in \boldsymbol{H}_M\},$$

we have

$$\ell_{\Phi_n}(\boldsymbol{z}) = \frac{\mathbb{G}\{\boldsymbol{A}_M\boldsymbol{z} + \boldsymbol{A}_M\boldsymbol{\Delta} + \boldsymbol{\omega} \in \boldsymbol{H}_M\}}{\Phi_n \times \mathbb{G}\{\boldsymbol{A}_M\boldsymbol{Z} + \boldsymbol{A}_M\boldsymbol{\Delta} + \boldsymbol{\omega} \in \boldsymbol{H}_M\}} = \frac{Q_1(\boldsymbol{z}; \boldsymbol{\Delta})}{\mathbb{E}_{\Phi_n}\left[Q_1(\boldsymbol{Z}; \boldsymbol{\Delta})\right]},$$

where we write the likelihood in terms of $\boldsymbol{z} = \sqrt{n}(\boldsymbol{t} - \boldsymbol{\mu})$. In this section, we take the randomization distribution to be Gaussian

$$\mathbb{G} = \mathcal{N}_d(\boldsymbol{0}, c\boldsymbol{\Sigma}_{\boldsymbol{\omega}}).$$

Assume that $\boldsymbol{\Sigma}_{\boldsymbol{\omega}} = c\boldsymbol{I}_d$ and that the selection region is rectangular $\boldsymbol{H}_M = \prod_{i=1}^n [b_i, \infty)$ for some vector $\boldsymbol{b} \in \mathbb{R}^d$.

Lemma 19 (Gaussian randomization: smoothness of the Gaussian likelihood) *Assume* (G) *holds. Then for any* $\kappa_2 < \frac{1}{2c}$*, there exists* κ_1 *such that*

$$e^{-\kappa_2\|\boldsymbol{A}_M\boldsymbol{z}\|_2^2} \cdot |\partial_{\boldsymbol{z}}^{\boldsymbol{\alpha}}\ell_{\Phi_n}(\boldsymbol{z})| = O\left(e^{\kappa_1 d(\boldsymbol{0}, \boldsymbol{H}_M - \boldsymbol{A}_M\boldsymbol{\Delta})^2}\right).$$

Proof

The denominator of the likelihood satisfies

$$\begin{aligned}
\mathbb{E}_{\Phi_n}\left[Q_1(\boldsymbol{Z}; \boldsymbol{\Delta})\right] &= \mathbb{P}_{\boldsymbol{Z}} \times \mathbb{P}_{\boldsymbol{\omega}}\{\boldsymbol{A}_M\boldsymbol{Z} + \boldsymbol{\omega} \in \boldsymbol{H}_M - \boldsymbol{A}_M\boldsymbol{\Delta}\} \\
&= \mathbb{P}_{\boldsymbol{\omega}' \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{A}_M\boldsymbol{\Sigma}\boldsymbol{A}_M^T + \boldsymbol{\Sigma}_{\boldsymbol{\omega}})}\{\boldsymbol{\omega}' \in \boldsymbol{H}_M - \boldsymbol{A}_M\boldsymbol{\Delta}\} \\
&= const \cdot \int_{\boldsymbol{H}_M} \exp\left(-\frac{1}{2}(\boldsymbol{\omega}' - \boldsymbol{A}_M\boldsymbol{\Delta})^T(\boldsymbol{A}_M\boldsymbol{\Sigma}\boldsymbol{A}_M^T + \boldsymbol{\Sigma}_{\boldsymbol{\omega}})^{-1}(\boldsymbol{\omega}' - \boldsymbol{A}_M\boldsymbol{\Delta})\right) d\boldsymbol{\omega}' \\
&= const \cdot \int_{\boldsymbol{H}_M} \exp\left(-\frac{1}{2}\left\|(\boldsymbol{A}_M\boldsymbol{\Sigma}\boldsymbol{A}_M^T + \boldsymbol{\Sigma}_{\boldsymbol{\omega}})^{-1/2}(\boldsymbol{\omega}' - \boldsymbol{A}_M\boldsymbol{\Delta})\right\|_2^2\right) d\boldsymbol{\omega}' \\
&\geqslant const \cdot \int_{\boldsymbol{H}_M} \exp\left(-\frac{1}{2}\|(\boldsymbol{A}_M\boldsymbol{\Sigma}\boldsymbol{A}_M^T + \boldsymbol{\Sigma}_{\boldsymbol{\omega}})^{-1/2}\|_2\left\|\boldsymbol{\omega}' - \boldsymbol{A}_M\boldsymbol{\Delta}\right\|_2^2\right) d\boldsymbol{\omega}' \\
&\geqslant const \cdot \exp\left(-\frac{\delta_1}{2}\|(\boldsymbol{A}_M\boldsymbol{\Sigma}\boldsymbol{A}_M^T + \boldsymbol{\Sigma}_{\boldsymbol{\omega}})^{-1/2}\|_2 \cdot d(\boldsymbol{0}, \boldsymbol{H}_M - \boldsymbol{A}_M\boldsymbol{\Delta})^2\right),
\end{aligned}$$

for some $\delta_1 > 1$.

The numerator of the pivot satisfies

$$\begin{aligned}|\partial_{\boldsymbol{z}}^{\boldsymbol{\alpha}} Q_1(\boldsymbol{z};\boldsymbol{\Delta})| &= const \cdot \left| \int\limits_{\boldsymbol{\omega}' \in \boldsymbol{H}_M} \partial_{\boldsymbol{z}}^{\boldsymbol{\alpha}} \exp\left(-\frac{1}{2c}\|\boldsymbol{\omega}' - \boldsymbol{A}_M\boldsymbol{z} - \boldsymbol{A}_M\boldsymbol{\Delta}\|_2^2\right) d\boldsymbol{\omega}' \right| \\ &\leqslant const \cdot \int\limits_{\boldsymbol{\omega}' \in \boldsymbol{H}_M} \exp\left(-\frac{\delta_2}{2c}\|\boldsymbol{\omega}' - \boldsymbol{A}_M\boldsymbol{z} - \boldsymbol{A}_M\boldsymbol{\Delta}\|_2^2\right) d\boldsymbol{\omega}' \\ &\leqslant const \cdot \exp\left(-\frac{\delta_2'}{2c} d\left(\boldsymbol{0}, \boldsymbol{H}_M - \boldsymbol{A}_M\boldsymbol{z} - \boldsymbol{A}_M\boldsymbol{\Delta}\right)^2\right),\end{aligned}$$

for some δ_2 and δ_2' satisfying $\delta_2' < \delta_2 < 1$.

By the triangle inequality

$$d(\boldsymbol{0}, \boldsymbol{H}_M - \boldsymbol{A}_M\boldsymbol{\Delta}) \leqslant d(\boldsymbol{0}, \boldsymbol{H}_M - \boldsymbol{A}_M\boldsymbol{z} - \boldsymbol{A}_M\boldsymbol{\Delta}) + \|\boldsymbol{A}_M\boldsymbol{z}\|_2,$$

hence

$$d(\boldsymbol{0}, \boldsymbol{H}_M - \boldsymbol{A}_M\boldsymbol{\Delta})^2 \leqslant 2d(\boldsymbol{0}, \boldsymbol{H}_M - \boldsymbol{A}_M\boldsymbol{z} - \boldsymbol{A}_M\boldsymbol{\Delta})^2 + 2\|\boldsymbol{A}_M\boldsymbol{z}\|_2^2.$$

Hence we get a further bound on the derivative of the numerator

$$|\partial_{\boldsymbol{z}}^{\boldsymbol{\alpha}} Q_1(\boldsymbol{z};\boldsymbol{\Delta})| \leqslant const \cdot \exp\left(-\frac{\delta_2'}{4c} d(\boldsymbol{0}, \boldsymbol{H}_M - \boldsymbol{A}_M\boldsymbol{\Delta})^2 + \frac{\delta_2'}{2c}\|\boldsymbol{A}_M\boldsymbol{z}\|_2^2\right).$$

Combining the upper bound on the derivative of the numerator together with the lower bound on the denominator we get that for any $\kappa_2 < \frac{1}{2c}$ ($\kappa_2 = \frac{\delta_2'}{2c}$), there exists κ_1 ($\kappa_1 = -\frac{\delta_2'}{4c} + \frac{\delta_1}{2}\|(\boldsymbol{A}_M\boldsymbol{\Sigma}\boldsymbol{A}_M^T + \boldsymbol{\Sigma}_{\boldsymbol{\omega}})^{-1/2}\|_2$) such that

$$|\partial_{\boldsymbol{z}}^{\boldsymbol{\alpha}} \ell_{\Phi_n}(\boldsymbol{z})| = \frac{|\partial_{\boldsymbol{z}}^{\boldsymbol{\alpha}} Q_1(\boldsymbol{z};\boldsymbol{\Delta})|}{\mathbb{E}_{\Phi_n}\left[Q_1(\boldsymbol{Z};\boldsymbol{\Delta})\right]} \leqslant const \cdot \exp\left(\kappa_1 d(\boldsymbol{0}, \boldsymbol{H}_M - \boldsymbol{A}_M\boldsymbol{\Delta})^2 + \kappa_2\|\boldsymbol{A}_M\boldsymbol{z}\|_2^2\right).$$

■

With Gaussian randomization, the pivot has a simper representation we now derive. Recall the variance of the randomization is $c\boldsymbol{I}_d$. Denote $\boldsymbol{v} = \boldsymbol{A}_M\boldsymbol{C}$, $\boldsymbol{w} = \boldsymbol{\omega}' - \boldsymbol{A}_M\boldsymbol{z}_A - \boldsymbol{A}_M\boldsymbol{\Delta}$. Assuming the matrix $\boldsymbol{A}$ is of size $d \times 1$, and $\boldsymbol{A}^T\boldsymbol{Z} =: T$ has variance 1, the numerator of the pivot multiplied by $(2\pi)^{(d+1)/2}c^{d/2}$ becomes:

$$\begin{aligned}&\int_{t \geqslant T} \exp\left(-\frac{t^2}{2}\right) \int_{\boldsymbol{H}_M} \exp\left(-\frac{1}{2c}(\boldsymbol{\omega}' - \boldsymbol{A}_M\boldsymbol{C}t - \boldsymbol{z}_A' - \boldsymbol{A}_M\boldsymbol{\Delta})^T(\boldsymbol{\omega}' - \boldsymbol{A}_M\boldsymbol{C}t - \boldsymbol{z}_A' - \boldsymbol{A}_M\boldsymbol{\Delta})\right) d\boldsymbol{\omega}' \\ &= \int_{t \geqslant T} \int_{\boldsymbol{H}_M} \exp\left(-\frac{1}{2}\left(\frac{1}{\|\boldsymbol{v}\|_2^2} + \frac{1}{c}\right)(\boldsymbol{v}t)^T(\boldsymbol{v}t) + \sqrt{\frac{1}{\|\boldsymbol{v}\|_2^2} + \frac{1}{c}}(\boldsymbol{v}t)^T\frac{1}{c}\sqrt{\frac{c\|\boldsymbol{v}\|_2^2}{\|\boldsymbol{v}\|_2^2 + c}}\boldsymbol{w} - \frac{1}{2c^2}\frac{\|\boldsymbol{v}\|_2^2 c}{\|\boldsymbol{v}\|_2^2 + c}\boldsymbol{w}^T\boldsymbol{w}\right) \\ &\qquad \cdot \exp\left(-\frac{1}{2}\boldsymbol{w}^T\boldsymbol{w}\left(\frac{1}{c} - \frac{\|\boldsymbol{v}\|_2^2}{c(\|\boldsymbol{v}\|_2^2 + c)}\right)\right) d\boldsymbol{\omega}' \\ &= \int_{\boldsymbol{H}_M} \int_{t \geqslant T} \exp\left(-\frac{1}{2}\left(\sqrt{\frac{\|\boldsymbol{v}\|_2^2 + c}{\|\boldsymbol{v}\|_2^2 c}}\boldsymbol{v}t - \frac{1}{c}\sqrt{\frac{\|\boldsymbol{v}\|_2^2 c}{\|\boldsymbol{v}\|_2^2 + c}}\boldsymbol{w}\right)^T\left(\sqrt{\frac{\|\boldsymbol{v}\|_2^2 + c}{\|\boldsymbol{v}\|_2^2 c}}\boldsymbol{v}t - \frac{1}{c}\sqrt{\frac{\|\boldsymbol{v}\|_2^2 c}{\|\boldsymbol{v}\|_2^2 + c}}\boldsymbol{w}\right)\right) dt \\ &\qquad \cdot \exp\left(-\frac{1}{2}\boldsymbol{w}^T\boldsymbol{w}\frac{1}{\|\boldsymbol{v}\|_2^2 + c}\right) d\boldsymbol{\omega}'.\end{aligned}$$

Denote

$$\tilde{\boldsymbol{v}} = \sqrt{\frac{\|\boldsymbol{v}\|_2^2 + c}{\|\boldsymbol{v}\|_2^2 c}}\boldsymbol{v}, \quad \tilde{\boldsymbol{w}} = \frac{1}{c}\sqrt{\frac{\|\boldsymbol{v}\|_2^2 c}{\|\boldsymbol{v}\|_2^2 + c}}\boldsymbol{w}.$$

Since

$$\begin{aligned}(\tilde{\boldsymbol{v}}t - \tilde{\boldsymbol{w}})^T(\tilde{\boldsymbol{v}}t - \tilde{\boldsymbol{w}}) &= \|\tilde{\boldsymbol{v}}\|_2^2\left(t^2 - 2t\frac{\tilde{\boldsymbol{v}}^T\tilde{\boldsymbol{w}}}{\|\tilde{\boldsymbol{v}}\|_2^2} + \frac{(\tilde{\boldsymbol{v}}^T\tilde{\boldsymbol{w}})^2}{\|\tilde{\boldsymbol{v}}\|_2^4}\right) - \frac{(\tilde{\boldsymbol{v}}^T\tilde{\boldsymbol{w}})2}{\|\tilde{\boldsymbol{v}}\|_2^2} + \|\tilde{\boldsymbol{w}}\|_2^2 \\ &= \|\tilde{\boldsymbol{v}}\|_2^2\left(t - \frac{\tilde{\boldsymbol{v}}^T\tilde{\boldsymbol{w}}}{\|\tilde{\boldsymbol{v}}\|_2^2}\right)^2 - \frac{(\tilde{\boldsymbol{v}}^T\tilde{\boldsymbol{w}})^2}{\|\tilde{\boldsymbol{v}}\|_2^2} + \|\tilde{\boldsymbol{w}}\|_2^2,\end{aligned}$$

we have the numerator (multiplied by $(2\pi)^{(d+1)/2}c^{d/2}$) to be

$$\begin{aligned}&\int_{\boldsymbol{H}_M}\int_{t\geqslant T}\exp\left(-\frac{\|\tilde{\boldsymbol{v}}\|_2^2}{2}\left(t - \frac{\tilde{\boldsymbol{v}}^T\tilde{\boldsymbol{w}}}{\|\tilde{\boldsymbol{v}}\|_2^2}\right)^2\right)dt \cdot \exp\left(\frac{(\tilde{\boldsymbol{v}}^T\tilde{\boldsymbol{w}})^2}{2\|\tilde{\boldsymbol{v}}\|_2^2} - \frac{1}{2}\|\tilde{\boldsymbol{w}}\|_2^2\right)\cdot\exp\left(-\frac{1}{2}\boldsymbol{w}^T\boldsymbol{w}\frac{1}{\|\boldsymbol{v}\|_2^2 + c}\right)d\boldsymbol{\omega}' \\ &= \frac{\sqrt{2\pi}}{\|\tilde{\boldsymbol{v}}\|_2}\int_{\boldsymbol{H}_M}\bar{\Phi}\left(\|\tilde{\boldsymbol{v}}\|_2\left(T - \frac{\tilde{\boldsymbol{v}}^T\tilde{\boldsymbol{w}}}{\|\tilde{\boldsymbol{v}}\|_2^2}\right)\right)\cdot\exp\left(\frac{(\tilde{\boldsymbol{v}}^T\tilde{\boldsymbol{w}})^2}{2\|\tilde{\boldsymbol{v}}\|_2^2} - \frac{1}{2}\|\tilde{\boldsymbol{w}}\|_2^2\right)\cdot\exp\left(-\frac{1}{2}\boldsymbol{w}^T\boldsymbol{w}\frac{1}{\|\boldsymbol{v}\|_2^2 + c}\right)d\boldsymbol{\omega}' \\ &= \frac{\sqrt{2\pi c}}{\sqrt{\|\boldsymbol{v}\|_2 + c}}\int_{\boldsymbol{H}_M}\bar{\Phi}\left(\sqrt{\frac{\|\boldsymbol{v}\|_2^2 + c}{c}}\left(T - \frac{\boldsymbol{v}^T\boldsymbol{w}}{\|\boldsymbol{v}\|_2^2 + c}\right)\right)\cdot\exp\left(\frac{(\boldsymbol{v}^T\boldsymbol{w})^2}{2c(\|\boldsymbol{v}\|_2^2 + c)} - \frac{1}{2c}\|\boldsymbol{w}\|_2^2\right)d\boldsymbol{\omega}' \\ &= \frac{\sqrt{2\pi c}}{\sqrt{\|\boldsymbol{v}\|_2 + c}}\int_{\boldsymbol{H}_M}\bar{\Phi}\left(\sqrt{\frac{\|\boldsymbol{v}\|_2^2 + c}{c}}\left(T - \frac{\boldsymbol{v}^T\boldsymbol{w}}{\|\boldsymbol{v}\|_2^2 + c}\right)\right)\cdot\exp\left(-\frac{1}{2}\boldsymbol{w}^T(\boldsymbol{v}\boldsymbol{v}^T + c\boldsymbol{I})^{-1}\boldsymbol{w}\right)d\boldsymbol{\omega}' \\ &= \frac{\sqrt{2\pi c}}{\sqrt{\|\boldsymbol{v}\|_2 + c}}\int_{\boldsymbol{H}_M}\bar{\Phi}\left(\sqrt{\frac{\|\boldsymbol{v}\|_2^2 + c}{c}}T - \frac{\boldsymbol{v}^T\boldsymbol{w}}{\sqrt{c(\|\boldsymbol{v}\|_2^2 + c)}}\right)\cdot\exp\left(-\frac{1}{2}\boldsymbol{w}^T(\boldsymbol{v}\boldsymbol{v}^T + c\boldsymbol{I})^{-1}\boldsymbol{w}\right)d\boldsymbol{\omega}',\end{aligned}$$

where $\bar{\Phi}(\cdot) = 1 - \Phi(\cdot)$ is the survival function of the standard Gaussian.

Hence the numerator becomes

$$\begin{aligned}&\frac{1}{\sqrt{|2\pi(\boldsymbol{v}\boldsymbol{v}^T + c\boldsymbol{I})|}}\int_{\boldsymbol{H}_M}\bar{\Phi}\left(\sqrt{\frac{\|\boldsymbol{v}\|_2^2 + c}{c}}T - \frac{\boldsymbol{v}^T\boldsymbol{w}}{\sqrt{c(\|\boldsymbol{v}\|_2^2 + c)}}\right)\cdot\exp\left(-\frac{1}{2}\boldsymbol{w}^T(\boldsymbol{v}\boldsymbol{v}^T + c\boldsymbol{I})^{-1}\boldsymbol{w}\right)d\boldsymbol{\omega}' \\ &= \int_{\boldsymbol{H}_M}\bar{\Phi}\left(\sqrt{\frac{\|\boldsymbol{v}\|_2^2 + c}{c}}T - \frac{\boldsymbol{v}^T\boldsymbol{w}}{\sqrt{c(\|\boldsymbol{v}\|_2^2 + c)}}\right)\cdot\phi_{(\boldsymbol{0},\boldsymbol{v}\boldsymbol{v}^T + \boldsymbol{\Sigma})}\left(\boldsymbol{\omega}' - \boldsymbol{A}_M\boldsymbol{z}_{\boldsymbol{A}} - \boldsymbol{A}_M\boldsymbol{\Delta}\right)d\boldsymbol{\omega}'.\end{aligned}$$

The denominator of the pivot is

$$\int_{\boldsymbol{H}_M}\phi_{(\boldsymbol{0},\boldsymbol{v}\boldsymbol{v}^T + \boldsymbol{\Sigma})}\left(\boldsymbol{\omega}' - \boldsymbol{A}_M\boldsymbol{z}_{\boldsymbol{A}} - \boldsymbol{A}_M\boldsymbol{\Delta}\right)d\boldsymbol{\omega}'.$$

The pivot becomes

$$\frac{\int_{\boldsymbol{H}_M}\bar{\Phi}\left(\sqrt{\frac{\|\boldsymbol{v}\|_2^2+c}{c}}T - \frac{\boldsymbol{v}^T\boldsymbol{w}}{\sqrt{c(\|\boldsymbol{v}\|_2^2+c)}}\right)\cdot\exp\left(-\frac{1}{2}\boldsymbol{w}^T(\boldsymbol{v}\boldsymbol{v}^T + c\boldsymbol{I})^{-1}\boldsymbol{w}\right)d\boldsymbol{\omega}'}{\int_{\boldsymbol{H}_M}\exp\left(-\frac{1}{2}\boldsymbol{w}^T(\boldsymbol{v}\boldsymbol{v}^T + c\boldsymbol{I})^{-1}\boldsymbol{w}\right)d\boldsymbol{\omega}'},$$

where $\boldsymbol{w} = \boldsymbol{\omega}' - \boldsymbol{A}_M\boldsymbol{z}_{\boldsymbol{A}} - \boldsymbol{A}_M\boldsymbol{\Delta}$.

Lemma 20 (Gaussian randomization: smoothness of the pivot) *Assume (G).*

$$|\partial_{\boldsymbol{z}}^{\boldsymbol{\alpha}}\mathcal{P}(\boldsymbol{z}, \boldsymbol{A}_M\boldsymbol{\Delta}, \boldsymbol{A}_M, \boldsymbol{H}_M)| \leqslant O(1 + d(\boldsymbol{0}, \boldsymbol{H}_M - \boldsymbol{A}_M\boldsymbol{\Delta})^3 + \|\boldsymbol{A}_M\boldsymbol{z}_{\boldsymbol{A}}\|_2^3)$$

Proof

The derivative of the pivot with respect to $\boldsymbol{z}$ is a linear combination of the terms:

$$\boldsymbol{J}_1 = \frac{\int_{\boldsymbol{H}_M} \left(\sqrt{\frac{\|\boldsymbol{v}\|_2^2+c}{c}}T - \frac{\boldsymbol{v}^T\boldsymbol{w}}{\sqrt{c(\|\boldsymbol{v}\|_2^2+c)}}\right) \phi\left(\sqrt{\frac{\|\boldsymbol{v}\|_2^2+c}{c}}T - \frac{\boldsymbol{v}^T\boldsymbol{w}}{\sqrt{c(\|\boldsymbol{v}\|_2^2+c)}}\right) \cdot \exp\left(-\frac{1}{2}\boldsymbol{w}^T(\boldsymbol{v}\boldsymbol{v}^T+c\boldsymbol{I})^{-1}\boldsymbol{w}\right) d\boldsymbol{\omega}'}{\int_{\boldsymbol{H}_M} \exp\left(-\frac{1}{2}\boldsymbol{w}^T(\boldsymbol{v}\boldsymbol{v}^T+c\boldsymbol{I})^{-1}\boldsymbol{w}\right) d\boldsymbol{\omega}'},$$

$$\boldsymbol{J}_2 = \frac{\int_{\boldsymbol{H}_M} \bar{\Phi}\left(\sqrt{\frac{\|\boldsymbol{v}\|_2^2+c}{c}}T - \frac{\boldsymbol{v}^T\boldsymbol{w}}{\sqrt{c(\|\boldsymbol{v}\|_2^2+c)}}\right) \cdot (\boldsymbol{v}\boldsymbol{v}^T+c\boldsymbol{I})^{-1}\boldsymbol{w} \cdot \exp\left(-\frac{1}{2}\boldsymbol{w}^T(\boldsymbol{v}\boldsymbol{v}^T+c\boldsymbol{I})^{-1}\boldsymbol{w}\right) d\boldsymbol{\omega}'}{\int_{\boldsymbol{H}_M} \exp\left(-\frac{1}{2}\boldsymbol{w}^T(\boldsymbol{v}\boldsymbol{v}^T+c\boldsymbol{I})^{-1}\boldsymbol{w}\right) d\boldsymbol{\omega}'}$$

and

$$\begin{aligned}\boldsymbol{J}_3 = &\frac{\int_{\boldsymbol{H}_M} \bar{\Phi}\left(\sqrt{\frac{\|\boldsymbol{v}\|_2^2+c}{c}}T - \frac{\boldsymbol{v}^T\boldsymbol{w}}{\sqrt{c(\|\boldsymbol{v}\|_2^2+c)}}\right) \cdot \exp\left(-\frac{1}{2}\boldsymbol{w}^T(\boldsymbol{v}\boldsymbol{v}^T+c\boldsymbol{I})^{-1}\boldsymbol{w}\right) d\boldsymbol{\omega}'}{\int_{\boldsymbol{H}_M} \exp\left(-\frac{1}{2}\boldsymbol{w}^T(\boldsymbol{v}\boldsymbol{v}^T+c\boldsymbol{I})^{-1}\boldsymbol{w}\right) d\boldsymbol{\omega}'} \\ &\cdot \frac{\int_{\boldsymbol{H}_M} (\boldsymbol{v}\boldsymbol{v}^T+c\boldsymbol{I})^{-1}\boldsymbol{w} \exp\left(-\frac{1}{2}\boldsymbol{w}^T(\boldsymbol{v}\boldsymbol{v}^T+c\boldsymbol{I})^{-1}\boldsymbol{w}\right) d\boldsymbol{\omega}'}{\int_{\boldsymbol{H}_M} \exp\left(-\frac{1}{2}\boldsymbol{w}^T(\boldsymbol{v}\boldsymbol{v}^T+c\boldsymbol{I})^{-1}\boldsymbol{w}\right) d\boldsymbol{\omega}'}.\end{aligned}$$

$\boldsymbol{J}_1$ is the expectation of a bounded function hence $\|\boldsymbol{J}_1\|_2$ is bounded. Since both $\|\boldsymbol{J}_2\|_2$ and $\|\boldsymbol{J}_3\|_2$ are upper bounded by

$$\|(\boldsymbol{v}\boldsymbol{v}^T+c\boldsymbol{I})^{-1}\|_2 \frac{\int_{\boldsymbol{H}_M} \|\boldsymbol{w}\|_2 \cdot \exp\left(-\frac{1}{2}\boldsymbol{w}^T(\boldsymbol{v}\boldsymbol{v}^T+c\boldsymbol{I})^{-1}\boldsymbol{w}\right) d\boldsymbol{\omega}'}{\int_{\boldsymbol{H}_M} \exp\left(-\frac{1}{2}\boldsymbol{w}^T(\boldsymbol{v}\boldsymbol{v}^T+c\boldsymbol{I})^{-1}\boldsymbol{w}\right) d\boldsymbol{\omega}'},$$

it suffices to analyze the growth of

$$\frac{\int_{\boldsymbol{H}_M} \|\boldsymbol{w}\|_2 \cdot \exp\left(-\frac{1}{2}\boldsymbol{w}^T(\boldsymbol{v}\boldsymbol{v}^T+c\boldsymbol{I})^{-1}\boldsymbol{w}\right) d\boldsymbol{\omega}'}{\int_{\boldsymbol{H}_M} \exp\left(-\frac{1}{2}\boldsymbol{w}^T(\boldsymbol{v}\boldsymbol{v}^T+c\boldsymbol{I})^{-1}\boldsymbol{w}\right) d\boldsymbol{\omega}'}$$

in order to see how fast the derivative of the pivot grows. Similarly, in order to bound the second and the third derivative of the pivot suffices to bound

$$\begin{aligned}&\frac{\int_{\boldsymbol{H}_M} \|\boldsymbol{w}\|_2^\alpha \cdot \exp\left(-\frac{1}{2}\boldsymbol{w}^T(\boldsymbol{v}\boldsymbol{v}^T+c\boldsymbol{I})^{-1}\boldsymbol{w}\right) d\boldsymbol{\omega}'}{\int_{\boldsymbol{H}_M} \exp\left(-\frac{1}{2}\boldsymbol{w}^T(\boldsymbol{v}\boldsymbol{v}^T+c\boldsymbol{I})^{-1}\boldsymbol{w}\right) d\boldsymbol{\omega}'} \\ &= \frac{\int_{\boldsymbol{H}_M - \boldsymbol{A}_M\boldsymbol{z}_{\boldsymbol{A}} - \boldsymbol{A}_M\boldsymbol{\Delta}} \|\boldsymbol{w}\|_2^\alpha \cdot \exp\left(-\frac{1}{2}\boldsymbol{w}^T(\boldsymbol{v}\boldsymbol{v}^T+c\boldsymbol{I})^{-1}\boldsymbol{w}\right) d\boldsymbol{w}}{\int_{\boldsymbol{H}_M - \boldsymbol{A}_M\boldsymbol{z}_{\boldsymbol{A}} - \boldsymbol{A}_M\boldsymbol{\Delta}} \exp\left(-\frac{1}{2}\boldsymbol{w}^T(\boldsymbol{v}\boldsymbol{v}^T+c\boldsymbol{I})^{-1}\boldsymbol{w}\right) d\boldsymbol{w}} \\ &= \mathbb{E}_{\boldsymbol{w}\sim\mathcal{N}(\boldsymbol{0},\boldsymbol{v}\boldsymbol{v}^T+c\boldsymbol{I})}\left[\|\boldsymbol{w}\|_2^\alpha \;\middle|\; \boldsymbol{w} \in \boldsymbol{H}_M - \boldsymbol{A}_M\boldsymbol{z}_{\boldsymbol{A}} - \boldsymbol{A}_M\boldsymbol{\Delta}\right]\end{aligned}$$

for $\alpha = 1, 2, 3$. Using Lemma 21, we get

$$|\partial_{\boldsymbol{z}}^{\boldsymbol{\alpha}} \mathcal{P}(\boldsymbol{z}, \boldsymbol{A}_M\boldsymbol{\Delta}, \boldsymbol{A}_M, \boldsymbol{H}_M)| \leqslant O(1 + d(\boldsymbol{0}, \boldsymbol{H}_M - \boldsymbol{A}_M\boldsymbol{z}_{\boldsymbol{A}} - \boldsymbol{A}_M\boldsymbol{\Delta})^3) \tag{65}$$

By the triangle inequality we get the conclusion.

■

The following lemma provides a bound on the growth of the moments of the Gaussian random variable after selection. Recall $\boldsymbol{H}_M = \prod_{i=1}^d [b_i, \infty)$ for some vector $\boldsymbol{b} \in \mathbb{R}^d$.

Lemma 21 *For a given $\boldsymbol{X} \sim \mathcal{N}_d(\mathbf{0}, \boldsymbol{\Sigma})$ and a vector $\boldsymbol{\mu} \in \mathbb{R}^d$, we have*

$$\mathbb{E}\left[\|\boldsymbol{X}\|_2^\alpha \mid \boldsymbol{X} \in \boldsymbol{H}_M - \boldsymbol{\mu}\right] = O(d(\mathbf{0}, \boldsymbol{H}_M - \boldsymbol{\mu})^\alpha + 1)$$

for $\alpha \in \mathbb{N}$.

Proof

We want to show $\mathbb{E}\left[\|\boldsymbol{X}\|_2^\alpha \mid \boldsymbol{X} \in \boldsymbol{H}_M - \boldsymbol{\mu}\right]$ grows at most linearly in $d(\mathbf{0}, \boldsymbol{H}_M - \boldsymbol{\mu})^\alpha$ for $d := d(\mathbf{0}, \boldsymbol{H}_M - \boldsymbol{\mu}) > 0$. The quantity above is increasing in d.

Denoting $\boldsymbol{Y} \sim \mathcal{N}\left(\mathbf{0}, \frac{1}{d^2}\boldsymbol{\Sigma}\right)$ we get the above equals

$$d^\alpha \cdot \mathbb{E}\left[\|\boldsymbol{Y}\|_2^\alpha \;\middle|\; \boldsymbol{Y} \in \frac{\boldsymbol{H}_M - \boldsymbol{\mu}}{d}\right].$$

Since $\boldsymbol{H}_M = \prod_{i=1}^d [b_i, \infty)$, $\frac{\boldsymbol{H}_M - \boldsymbol{\mu}}{d}$ is a rectangle $\prod_{i=1}^d [u_i, \infty)$, $\|\boldsymbol{u}\|_2 = 1$. So $\mathbb{E}\left[\|\boldsymbol{Y}\|_2^\alpha \;\middle|\; \boldsymbol{Y} \in \frac{\boldsymbol{H}_M - \boldsymbol{\mu}}{d}\right]$ is decreasing in d hence bounded above.

■

C.2. Applying Chatterjee's theorem for Gaussian randomization

Lemma 22 (Gaussian randomization: closeness of Gaussian and non-parametric LR) *Assuming* (AL), (LA-w), (NA) *and* (G) *and* (subG) *hold, we have*

$$\mathbb{E}_{\mathbb{F}_n}\left[|\ell_{\mathbb{F}_n}(\boldsymbol{Z}) - \ell_{\Phi_n}(\boldsymbol{Z})|\right] \leqslant \frac{1}{\sqrt{n}} C(C_M, K_g, K_G, \boldsymbol{\Sigma}),$$

where the constant on the RHS depends only on its arguments.

Proof

Since

$$\begin{aligned}
\mathbb{E}_{\mathbb{F}_n}\left[|\ell_{\mathbb{F}_n}(\boldsymbol{Z}) - \ell_{\Phi_n}(\boldsymbol{Z})|\right] &= \mathbb{E}_{\mathbb{F}_n}\left[Q_1(\boldsymbol{Z}; \boldsymbol{\Delta})\right] \left| \frac{1}{\mathbb{E}_{\mathbb{F}_n}\left[Q_1(\boldsymbol{Z}; \boldsymbol{\Delta})\right]} - \frac{1}{\mathbb{E}_{\Phi_n}\left[Q_1(\boldsymbol{Z}; \boldsymbol{\Delta})\right]} \right| \\
&= \frac{\left|\mathbb{E}_{\mathbb{F}_n}\left[Q_1(\boldsymbol{Z}; \boldsymbol{\Delta})\right] - \mathbb{E}_{\Phi_n}\left[Q_1(\boldsymbol{Z}; \boldsymbol{\Delta})\right]\right|}{\mathbb{E}_{\Phi_n}\left[Q_1(\boldsymbol{Z}; \boldsymbol{\Delta})\right]} = \left|\mathbb{E}_{\mathbb{F}_n}[\ell_{\Phi_n}(\boldsymbol{Z})] - \mathbb{E}_{\Phi_n}[\ell_{\Phi_n}(\boldsymbol{Z})]\right|.
\end{aligned}$$

We use Chatterjee technique here, providing an upper bound on the quantity above using the smoothness of $\ell_{\Phi_n}(\boldsymbol{z})$. From Lemma 19, we have

$$|\partial_{\boldsymbol{z}}^{\boldsymbol{\alpha}} \ell_{\Phi_n}(\boldsymbol{z})| \leqslant const \cdot \exp\left(\kappa_1 d(\mathbf{0}, \boldsymbol{H}_M - \boldsymbol{A}_M \boldsymbol{\Delta})^2 + \kappa_2 \|\boldsymbol{A}_M \boldsymbol{z}\|_2^2\right),$$

implying

$$\lambda_3^{(n)}(\ell_{\Phi_n}) \leqslant \frac{1}{n^{3/2}} \exp(\kappa_1 d(\mathbf{0}, \boldsymbol{H}_M - \boldsymbol{A}_M \boldsymbol{\Delta})^2).$$

for $\Omega_n(\boldsymbol{z}) = \|\boldsymbol{A}_M \boldsymbol{z}\|_2^2$, $\boldsymbol{z} \in \mathbb{R}^p$, and $\mathcal{W}_n(x) = e^{\kappa_2 x}$, $x \in \mathbb{R}$. Using Theorem 16, we have

$$\begin{aligned}
&\left|\mathbb{E}_{\mathbb{F}_n}\left[\ell_{\Phi_n}(\boldsymbol{Z})\right] - \mathbb{E}_{\Phi_n}\left[\ell_{\Phi_n}(\boldsymbol{Z})\right]\right| \\
&\leqslant \frac{1}{2}\sum_{i=1}^n \lambda_3^{(n)}(\ell_{\Phi_n}) \mathbb{E}\left[e^{2\kappa_2 \|\boldsymbol{A}_M \mathcal{M}(\boldsymbol{W}_i)\|_2^2}\right] \left(\mathbb{E}\left[\|\boldsymbol{g}_i\|_1^3\right] + \mathbb{E}\left[\|\boldsymbol{y}_i\|_1^3\right]\right) \\
&+ \frac{1}{2}\sum_{i=1}^n \lambda_3^{(n)}(\ell_{\Phi_n}) \mathbb{E}\left[e^{2\kappa_2 \|\boldsymbol{A}_M \mathcal{M}(\widetilde{\boldsymbol{W}}_i - \boldsymbol{W}_i)\|_2^2} \|\boldsymbol{g}_i\|_1^3\right] \\
&+ \frac{1}{2}\sum_{i=1}^n \lambda_3^{(n)}(\ell_{\Phi_n}) \mathbb{E}\left[e^{2\kappa_2 \|\boldsymbol{A}_M \mathcal{M}(\widetilde{\boldsymbol{W}}_{i-1} - \boldsymbol{W}_i)\|_2^2} \|\boldsymbol{y}_i\|_1^3\right]
\end{aligned}$$

By the sub-Gaussain assumption and the local alternatives the sum in the RHS above converges to zero as n tends to infinity. ■

Theorem 23 *Assume* (AL), (LA-w), (NA) *and* (G) *and* (subG) *hold. Given a function* $\mathcal{H}:\mathbb{R}\to\mathbb{R}$ *with uniformly bounded derivatives up to the third order (call this bound* $K_{\mathcal{H}}$*), we have*

$$\left|\mathbb{E}_{\mathbb{F}_n^*}\mathcal{H}(\mathcal{P}_n(\boldsymbol{Z}))-\mathbb{E}_{\Phi_n^*}\mathcal{H}(\mathcal{P}_n(\boldsymbol{G}))\right|=o_n(1).$$

Proof Similar to the proof of the above Lemma. ■

D. Bootstrap consistency - additional lemmas

The following lemma is an important step in proving asymptotics of $\mathcal{P}^*$. It will be used to show that the denominator of the test statistic $\mathcal{P}^*$ is bounded below in probability uniformly over $\mathbb{F}_n\in\mathcal{F}_n$. Without randomization, the result of the lemma would not hold.

Lemma 24 (Lower bound) *Assuming* (CLT), (LA), (Lip) *and* (NA), *we have that for every* $\delta>0$, *there exists* $C>0$ *such that*

$$\sup_{\mathbb{F}_n\in\mathcal{F}_n}\mathbb{F}_n\left\{Q_2(\boldsymbol{Z}_{\boldsymbol{A}};\boldsymbol{\Delta})<C\right\}<\delta$$

for n *sufficiently large.*

Proof

Using the Lipschitz assumption on the randomization, we have

$$\begin{aligned}
Q_2(\boldsymbol{\tau};\boldsymbol{z}_{\boldsymbol{A}},\boldsymbol{\Delta}) &= \frac{1}{C_g}\int_{\boldsymbol{H}_M-\boldsymbol{A}_M\boldsymbol{\Delta}}\exp\left(-\tilde{g}(\boldsymbol{\omega}'-\boldsymbol{A}_M\boldsymbol{z}_{\boldsymbol{A}}-\boldsymbol{A}_M\boldsymbol{C}\boldsymbol{\tau})\right)d\boldsymbol{\omega}'\\
&= \frac{1}{C_g}\int_{\boldsymbol{H}_M-\boldsymbol{A}_M\boldsymbol{\Delta}}\exp\left(-\tilde{g}(\boldsymbol{\omega}'-\boldsymbol{A}_M\boldsymbol{\Delta})-K_g\|\boldsymbol{A}_M\boldsymbol{z}_{\boldsymbol{A}}\|_h-K_g\|\boldsymbol{A}_M\boldsymbol{C}\boldsymbol{\tau}\|_h\right)d\boldsymbol{\omega}'\\
&= \exp\left(-K_g\|\boldsymbol{A}_M\boldsymbol{z}_{\boldsymbol{A}}\|_h\right)\exp\left(-K_g\|\boldsymbol{A}_M\boldsymbol{C}\boldsymbol{\tau}\|_h\right)\int_{\boldsymbol{H}_M-\boldsymbol{A}_M\boldsymbol{\Delta}}\exp(-\tilde{g}(\boldsymbol{\omega}'))d\boldsymbol{\omega}'
\end{aligned}$$

for all $\boldsymbol{\tau}\in\mathbb{R}^a$ and $\boldsymbol{z}_{\boldsymbol{A}}\in\mathbb{R}^p$. Since $d_h\left(\boldsymbol{0},\boldsymbol{H}_M-\boldsymbol{A}_M\boldsymbol{\Delta}\right)\leqslant B$ (local alternatives), the last integral is O(1).

$$Q_2\left(\boldsymbol{\tau};\boldsymbol{z}_{\boldsymbol{A}},\boldsymbol{\Delta}\right)\geqslant e^{-K_gC_M\|\boldsymbol{z}_{\boldsymbol{A}}\|_2}e^{-K_gC_M\|\boldsymbol{C}\boldsymbol{\tau}\|_2}\int_{\boldsymbol{H}_M-\boldsymbol{A}_M\boldsymbol{\Delta}}\exp(-\tilde{g}(\boldsymbol{\omega}'))d\boldsymbol{\omega}',$$

for all $\boldsymbol{\tau}\in\mathbb{R}^a$ and $\boldsymbol{z}_{\boldsymbol{A}}\in\mathbb{R}^p$, where the second inequality follows from (NA) assumption for n large enough. Defining a function of $x\in\mathbb{R}$

$$\int_{\mathbb{R}^a}e^{-\|\boldsymbol{C}\boldsymbol{\tau}\|_2x}\frac{1}{(2\pi)^{-a/2}|\boldsymbol{\Sigma}_{\boldsymbol{A}}|^{-1/2}}e^{-\frac{1}{2}\boldsymbol{\tau}\boldsymbol{\Sigma}_{\boldsymbol{A}}^{-1}\boldsymbol{\tau}}d\boldsymbol{\tau}=:f_{\boldsymbol{\Sigma}}(x),$$ [23]

we have

$$Q_3\left(\boldsymbol{z}_{\boldsymbol{A}};\boldsymbol{\Delta}\right)\geqslant e^{-K_gC_M\|\boldsymbol{z}_{\boldsymbol{A}}\|_2}f_{\boldsymbol{\Sigma}}(K_gC_M)\int_{\boldsymbol{H}_M-\boldsymbol{A}_M\boldsymbol{\Delta}}\exp(-\tilde{g}(\boldsymbol{\omega}'))d\boldsymbol{\omega}'$$

[23] Since under $\boldsymbol{\tau}\sim\mathcal{N}_a(\boldsymbol{0},\boldsymbol{\Sigma}_{\boldsymbol{A}})$, $\boldsymbol{C}\boldsymbol{\tau}\sim\mathcal{N}_p(\boldsymbol{0},\boldsymbol{\Sigma})$, hence the RHS does not depend on $\boldsymbol{A}$.

for all $\boldsymbol{z_A} \in \mathbb{R}^d$.

For a given $\delta > 0$, it suffices to find a constant $C_1 > 0$ such that

$$\sup_{\mathbb{F}_n \in \mathcal{F}} \mathbb{F}_n \left\{ \|\boldsymbol{Z_A}\|_2 > C_1 \right\} < \delta$$

for n large enough. By (CLT) assumption, $\boldsymbol{Z_A}$ converges uniformly in distribution to a random variable $\boldsymbol{G} \sim \mathcal{N}_p\left(\boldsymbol{0}, (\boldsymbol{I}_p - \boldsymbol{C}\boldsymbol{A}^T)\boldsymbol{\Sigma}(\boldsymbol{I}_p - \boldsymbol{C}\boldsymbol{A}^T)^T\right)$. Using the uniform continuous mapping theorem we get that $\|\boldsymbol{Z_A}\|_2$ converges uniformly in distribution to $\|\boldsymbol{G}\|_2$, i.e.

$$\lim_{n\to\infty} \sup_{\mathbb{F}_n \in \mathcal{F}_n} \sup_{t \in \mathbb{R}} |\mathbb{F}_n \left\{ \|\boldsymbol{Z_A}\|_2 \leqslant t \right\} - \mathbb{P}_{\boldsymbol{G}} \left\{ \|\boldsymbol{G}\|_2 \leqslant t \right\}| = 0.$$

This implies (D) holds.

■

The following lemma justifies using $\hat{\boldsymbol{T}}_{\boldsymbol{A}}$ instead of $\boldsymbol{T_A}$ in constructing the bootstrap pivot

Lemma 25 *Assuming* (CLT), (NA) *and* (Var) *hold, for any* $\delta > 0$ *the following holds*

$$\lim_{n\to\infty} \sup_{\mathbb{F}_n \in \mathcal{F}_n} \mathbb{F}_n \left\{ \left\| \sqrt{n}\boldsymbol{A}_M(\hat{\boldsymbol{T}}_{\boldsymbol{A}} + \hat{\boldsymbol{C}}\boldsymbol{A}^T\boldsymbol{\mu}) - \sqrt{n}\boldsymbol{A}_M\left(\boldsymbol{T_A} + \boldsymbol{C}\boldsymbol{A}^T\boldsymbol{\mu}\right) \right\|_2 > \delta \right\} = 0.$$

Proof Since

$$\sqrt{n}\boldsymbol{A}_M(\hat{\boldsymbol{T}}_{\boldsymbol{A}} - \boldsymbol{T_A}) + \sqrt{n}\boldsymbol{A}_M(\hat{\boldsymbol{C}} - \boldsymbol{C})\boldsymbol{A}^T\boldsymbol{\mu} = \boldsymbol{A}_M(\boldsymbol{C} - \hat{\boldsymbol{C}})\left(\sqrt{n}\left(\boldsymbol{A}^T\boldsymbol{T} - \boldsymbol{A}^T\boldsymbol{\mu}\right)\right),$$

we have

$$\begin{aligned} &\left\| \sqrt{n}\boldsymbol{A}_M(\hat{\boldsymbol{T}}_{\boldsymbol{A}} + \hat{\boldsymbol{C}}\boldsymbol{A}^T\boldsymbol{\mu}) - \sqrt{n}\boldsymbol{A}_M(\boldsymbol{T_A} + \boldsymbol{C}\boldsymbol{A}^T\boldsymbol{\mu}) \right\|_2 \\ &\leqslant \|\boldsymbol{A}_M\|_2 \|\boldsymbol{C} - \hat{\boldsymbol{C}}\|_2 \left\| \sqrt{n}\left(\boldsymbol{A}^T(\boldsymbol{T} - \boldsymbol{\mu})\right) \right\|_2 \\ &\leqslant a_M \|\boldsymbol{C} - \hat{\boldsymbol{C}}\|_2 \left\| \sqrt{n}(\boldsymbol{A}^T\boldsymbol{T} - \boldsymbol{A}^T\boldsymbol{\mu}) \right\|_2, \end{aligned} \tag{66}$$

where the second inequality holds for n large enough by (NA) assumption. Since $\|\hat{\boldsymbol{C}} - \boldsymbol{C}\|_2 = o_{\mathbb{F}_n}(1)$ uniformly across $\mathbb{F}_n$ by assumption (Var) and $\left\| \sqrt{n}(\boldsymbol{A}^T\boldsymbol{T} - \boldsymbol{A}^T\boldsymbol{\mu}) \right\|_2$ converges in distribution by (CLT) assumption to $\|\boldsymbol{G}\|_2$, where $\boldsymbol{G} \sim \mathcal{N}_a\left(\boldsymbol{0}, \boldsymbol{\Sigma_A}\right)$ uniformly across $\mathbb{F}_n \in \mathcal{F}_n$, we have that the term in (66) converges to zero in probability uniformly across $\mathbb{F}_n \in \mathcal{F}$, hence the conclusion follows. ■

E. Model selection and the asymptotics for the Lasso

Here we characterize the model $(E, \boldsymbol{s}_E)$ chosen by the randomized Lasso objective in Section 5.1 and show that asymptotically the selection region is affine in terms of a test statistic we introduce. The KKT conditions for the randomized LASSO are

$$\begin{aligned} \boldsymbol{X}_E^T\left(\boldsymbol{y} - \boldsymbol{X}_E\hat{\boldsymbol{\beta}}_E\right) &= \lambda \boldsymbol{s}_E - \boldsymbol{\omega}_E + \epsilon\hat{\boldsymbol{\beta}}_E \\ \boldsymbol{X}_{-E}^T\left(\boldsymbol{y} - \boldsymbol{X}_E\hat{\boldsymbol{\beta}}_E\right) &= \lambda \boldsymbol{s}_{-E} - \boldsymbol{\omega}_{-E} \\ \operatorname{diag}(\boldsymbol{s}_E)\hat{\boldsymbol{\beta}}_E \geqslant 0,\ \ \|\boldsymbol{s}_{-E}\|_\infty &\leqslant 1, \end{aligned} \tag{67}$$

where $\boldsymbol{s}_{-E}$ is the sub-gradient of $\frac{\partial|\beta_j|}{\partial\beta_j}\big|_{\boldsymbol{\beta}=\hat{\boldsymbol{\beta}}}$ for $j \notin E$ (inactive variables) and $\operatorname{diag}(\boldsymbol{s}_E)$ is an $|E| \times |E|$ matrix having the entries of $\boldsymbol{s}_E$ on the diagonal and zeros elsewhere. Let us denote $\bar{\boldsymbol{\beta}}_E =$

$\left(\boldsymbol{X}_E^T\boldsymbol{X}_E\right)^{-1}\boldsymbol{X}_E^T\boldsymbol{y}$, the least squares solution using $\boldsymbol{y}$ as the response and $\boldsymbol{X}_E$ as the matrix of predictors. Then the KKT conditions above become

$$\begin{aligned}\widehat{\boldsymbol{\beta}}_E &= \bar{\boldsymbol{\beta}}_E - (\boldsymbol{X}_E^T\boldsymbol{X}_E)^{-1}(\lambda \boldsymbol{s}_E - \boldsymbol{\omega}_E) + \epsilon(\boldsymbol{X}_E^T\boldsymbol{X}_E)^{-1}\widehat{\boldsymbol{\beta}}_E \\ \lambda \boldsymbol{s}_{-E} &= \boldsymbol{X}_{-E}^T\left(\boldsymbol{y} - \boldsymbol{X}_E\bar{\boldsymbol{\beta}}_E\right) + \boldsymbol{X}_{-E}^T\boldsymbol{X}_E(\boldsymbol{X}_E^T\boldsymbol{X}_E)^{-1}(\lambda\boldsymbol{s}_E - \boldsymbol{\omega}_E) \\ &\quad + \epsilon\boldsymbol{X}_{-E}^T\boldsymbol{X}_E(\boldsymbol{X}_E^T\boldsymbol{X}_E)^{-1}\widehat{\boldsymbol{\beta}}_E + \boldsymbol{\omega}_{-E}.\end{aligned}$$

Since by the strong law of large numbers

$$\begin{aligned}\boldsymbol{X}_E^T\boldsymbol{X}_E &= \sum_{i=1}^n \boldsymbol{x}_{E,i}\boldsymbol{x}_{E,i}^T \xrightarrow{\mathbb{F}_n} \mathbb{E}_{\mathbb{F}_n}\left[n\cdot\boldsymbol{x}_{E,1}\boldsymbol{x}_{E,1}^T\right] =: \boldsymbol{E}_1 \in \mathbb{R}^{|E|\times|E|} \\ \boldsymbol{X}_{-E}^T\boldsymbol{X}_E &= \sum_{i=1}^n \boldsymbol{x}_{-E,i}\boldsymbol{x}_{E,i}^T \xrightarrow{\mathbb{F}_n} \mathbb{E}_{\mathbb{F}_n}\left[n\cdot\boldsymbol{x}_{-E,1}\boldsymbol{x}_{E,1}^T\right] =: \boldsymbol{E}_2 \in \mathbb{R}^{(p-|E|)\times|E|},\end{aligned}$$

where the deviations from the mean are $O_{\mathbb{F}_n}\left(\frac{1}{\sqrt{n}}\right)$ (under CLT assumptions), we write the inequalities from (67) using

$$\begin{aligned}&\operatorname{diag}(\boldsymbol{s}_E)\bar{\boldsymbol{\beta}}_E - \operatorname{diag}(\boldsymbol{s}_E)\boldsymbol{E}_1^{-1}(\lambda\boldsymbol{s}_E - \boldsymbol{\omega}_M) + O_{\mathbb{F}_n}\left(\frac{1}{\sqrt{n}}\right) \geqslant 0, \\ &-\lambda \leqslant \left\|\boldsymbol{X}_{-E}^T(\boldsymbol{y} - \boldsymbol{X}_E\bar{\boldsymbol{\beta}}_E) + \boldsymbol{E}_2\boldsymbol{E}_1^{-1}(\lambda\boldsymbol{s}_E - \boldsymbol{\omega}_M) + \boldsymbol{\omega}_{-E} + O_{\mathbb{F}_n}\left(\frac{1}{\sqrt{n}}\right)\right\|_\infty \leqslant \lambda.\end{aligned}$$

hence the selection event is asymptotically affine in terms of the following test statistic

$$\boldsymbol{D} = \begin{pmatrix}\boldsymbol{D}_E \\ \boldsymbol{D}_{-E}\end{pmatrix} = \begin{pmatrix}\bar{\boldsymbol{\beta}}_E \\ \boldsymbol{X}_{-E}^T\left(\boldsymbol{y} - \boldsymbol{X}_E\bar{\boldsymbol{\beta}}_E\right)\end{pmatrix}$$

The next lemma shows that $\frac{1}{\sqrt{n}}\boldsymbol{D}$ is asymptotically linear as defined in [Tian and Taylor, 2015b], i.e. $\boldsymbol{D} = \frac{1}{\sqrt{n}}\sum_{i=1}^n \boldsymbol{\xi}_i + o_{\mathbb{F}_n}(1)$, where $\boldsymbol{\xi}_i$ are measurable with respect to $(\boldsymbol{X}_i, y_i)$ and $\mathbb{E}_{\mathbb{F}_n}\boldsymbol{\xi}_i$ is the same across $i = 1, \ldots, n$.

Lemma 26 (Asymptotic Linearity) *Test statistic $\frac{1}{\sqrt{n}}\boldsymbol{D}$ is asymptotically linear.*

Proof We have

$$\begin{aligned}&\bar{\boldsymbol{\beta}}_E - \boldsymbol{\beta}_E^* = (\boldsymbol{X}_E^T\boldsymbol{X}_E)^{-1}\boldsymbol{X}_E^T(\boldsymbol{y} - \boldsymbol{X}_E\boldsymbol{\beta}_E^*) \\ &= \boldsymbol{E}_1^{-1}\boldsymbol{X}_E^T(\boldsymbol{y} - \boldsymbol{X}_E\boldsymbol{\beta}_E^*) + \left((\boldsymbol{X}_E^T\boldsymbol{X}_E)^{-1} - \boldsymbol{E}_1^{-1}\right)\boldsymbol{X}_E^T(\boldsymbol{y} - \boldsymbol{X}_E\boldsymbol{\beta}_E^*) \\ &= \boldsymbol{E}_1^{-1}\sum_{i=1}^n\left(y_i - \boldsymbol{x}_{E,i}^T\boldsymbol{\beta}_E^*\right)\boldsymbol{x}_{E,i} + O_{\mathbb{F}_n}\left(\frac{1}{\sqrt{n}}\right)O_{\mathbb{F}_n}(1),\end{aligned}$$

and the last term is $O_{\mathbb{F}_n}\left(\frac{1}{\sqrt{n}}\right)$, which implies $o_{\mathbb{F}_n}(1)$. The second term similarly satisfies

$$\begin{aligned}&\boldsymbol{X}_{-E}^T\left(\boldsymbol{y} - \boldsymbol{X}_E\bar{\boldsymbol{\beta}}_E\right) - \boldsymbol{X}_{-E}^T(\boldsymbol{y} - \boldsymbol{X}_E\boldsymbol{\beta}_E^*) = \boldsymbol{X}_{-E}^T\boldsymbol{X}_E(\boldsymbol{\beta}_E^* - \bar{\boldsymbol{\beta}}_E) \\ &= -\boldsymbol{X}_{-E}^T\boldsymbol{X}_E(\boldsymbol{X}_E^T\boldsymbol{X}_E)^{-1}\boldsymbol{X}_E^T(\boldsymbol{y} - \boldsymbol{X}_E\boldsymbol{\beta}_E^*) \\ &= -\boldsymbol{E}_2\boldsymbol{E}_1^{-1}\boldsymbol{X}_E^T(\boldsymbol{y} - \boldsymbol{X}_E\boldsymbol{\beta}_E^*) - \left(\boldsymbol{X}_{-E}^T\boldsymbol{X}_E(\boldsymbol{X}_E^T\boldsymbol{X}_E)^{-1} - \boldsymbol{E}_2\boldsymbol{E}_1^{-1}\right)\boldsymbol{X}_E^T(\boldsymbol{y} - \boldsymbol{X}_E\boldsymbol{\beta}_E).\end{aligned}$$

The first term after the last equality is linear and the second term is of order $O_{\mathbb{F}_n}\left(\frac{1}{\sqrt{n}}\right)O_{\mathbb{F}_n}(1)$, hence also $o_{\mathbb{F}_n}(1)$.

■